APPLICATION AND ANALYSIS OF COGNITIVE SEARCH MODELS ON COLLABORATIVE MEMORY TASKS

Willa M. Mannering

Submitted to the faculty of the University Graduate School in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Psychological and Brain Sciences and the Program in Cognitive Science,

Indiana University

June 2023

Accepted by the Graduate Faculty of Indiana University in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Doctoral Committee

Michael N. Jones, Ph.D.

Richard M. Shiffrin, Ph.D.

Peter M. Todd, Ph.D.

Sandra Kübler, Ph.D.

May 11, 2023

Acknowledgements

I would like to express my sincere gratitude to everyone who has supported me throughout my journey in pursuing a doctoral degree in psychology and cognitive science at Indiana University.

First and foremost, I would like to thank my advisor, Michael Jones, for his guidance, support, and patience throughout my graduate studies. His expertise, encouragement, and constructive feedback have been invaluable to my growth as a researcher and scholar. I would also like to extend my gratitude to my committee members, Rich Shiffrin, Peter Todd, Sandra Kübler, and Rob Goldstone for their time, insight, and invaluable feedback. Their diverse perspectives and expertise have greatly enriched my work and research. I would like to thank both Peter and Rob for their involvement in my undergraduate career as educators and mentors. Without my experiences with them in the classroom and the lab I would not be the researcher I am today. I am grateful to my collaborator, Suparna Rajaram, for her insight and participation in several studies completed for this dissertation.

I would also like to thank Thomas Busey for his invaluable guidance and support in training me as an undergraduate researcher. My academic career as it is now truly began in his lab where he provided me with numerous opportunities to engage in hands-on research experiences. Through my tenure in his lab I developed critical thinking, data analysis, and communication skills and was able to disseminate my research. His mentorship as an undergraduate prepared me for the work I would face as a graduate student.

In addition, I am grateful to my fellow IU graduate students, collaborators, and friends who have provided me with support, camaraderie, and intellectual stimulation throughout my studies. Their friendship, encouragement, and collaboration have made this journey possible. I will always remember these 5 years in Bloomington fondly because of our friendship. Beyond friends made in graduate school, I would also like to thank older friends from my undergraduate studies and earlier. You have been an invaluable support system not only for matters related to my degree but personal matters as well.

Lastly, I would like to thank my family. My parents, Fred and Jill, and my sister, Freyda, have provided unwavering support and encouragement, which have sustained me during the most challenging times of my studies. My parents always encouraged me to pursue higher education and their support was pivotal in my decision to pursue a doctoral degree. Their steadfast belief in me and my abilities has been a constant source of motivation and inspiration throughout my entire life. Thank you to my partner, Josh Apanavicius, for being amazing and supportive as I've pursued my education. Finally, thank you to Dr. Oswald Paws for being a fluffy, cuddly cat and providing much needed emotional support.

Willa M. Mannering

APPLICATION AND ANALYSIS OF COGNITIVE SEARCH MODELS ON COLLABORATIVE MEMORY TASKS

While humans often encode and retrieve memories in groups, the bulk of our knowledge of human memory comes from paradigms with individuals in isolation. The primary phenomenon of interest within the field of collaborative memory is *collaborative inhibition*: the tendency for collaborative groups to underperform in free recall tasks compared to nominal groups of the same size. Most research in this field is led by verbal conceptual theories without guidance from formal computational models. The goal of this dissertation was to expand the modeling efforts for collaborative memory within both episodic and semantic memory tasks by developing a formal computational model of collaborative recall. To this aim, a collaborative framework to scale the Search of Associative Memory model (SAM; Raaijmakers & Shiffrin, 1981) to collaborative free recall paradigms (dubbed cSAM) with multiple models working together was presented. This work shows that SAM, adapted to cSAM, can act as a unified theory to explain both individual and collaborative memory effects and can also provide insight into mechanistic explanations of collaborative inhibition that would be difficult or nearly impossible to study behaviorally.

Additionally, collaborative inhibition in semantic memory tasks has received concerningly little attention from the field and existing studies (of which there are only 2) have employed a task more similar to fact retrieval paradigms (Andersson & Ronnberg, 1996; Weldon, 2000). To remedy this gap in knowledge, a novel collaborative verbal fluency task was employed to determine whether collaborative inhibition is present in semantic memory tasks in which category structure is prominent. Furthermore, individual and group search behaviors were analyzed using an optimal foraging model of semantic memory. Collaborative inhibition was present in the

collaborative verbal fluency task and cognitive search behavior differed significantly between collaborative and nominal groups. Altogether, this work significantly contributes to the modeling and experimental efforts within the collaborative memory field which will help the field generate new predictions and experiments to advance the study of memory storage and retrieval.

Michael N. Jones, Ph.D.

Richard M. Shiffrin, Ph.D.

Peter M. Todd, Ph.D.

Sandra Kübler, Ph.D.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

Collaborative memory is a relatively new field of study, first gaining attention from cognitive scientists in the late 1990s, which focuses on group interaction and remembering. Whereas previous researchers of group memory took a social and historical approach, collaborative memory researchers take a distinctly cognitive and empirical approach to studying group memory. Research in collaborative memory is focused on the internal cognitive mechanisms responsible for group memory phenomena and takes inspiration from experimental paradigms and theories originating from the individual memory field. The primary focus within the field is on an effect called *collaborative inhibition*—the tendency for collaborative groups to underperform in recall tasks compared to nominal groups of the same size.

Currently, there are several viable mechanistic explanations for collaborative inhibition, but the explanation with the most empirical support is the retrieval disruption hypothesis. This hypothesis posits that the inhibitory effect of collaboration occurs because individual retrieval strategies are disrupted during group activities (B. H. Basden, Basden, Bryner, & Thomas, 1997).

According to this hypothesis, each group member develops an idiosyncratic method of organizing information in memory during the study phase of recall which is then disrupted by incongruous cues produced by other group members during the test phase of recall. While a large body of experimental research exists within the collaborative memory field, primarily investigating the predictions of the retrieval disruption hypothesis on episodic free recall, there is a distinct lack of modeling efforts and research related to semantic memory tasks. Given the tendency of the collaborative memory field to borrow and test mechanistic theories originally developed within the field of individual memory, the same approach should be taken to create a cognitive model of collaborative memory. The introduction of this dissertation will present a discussion of the history of the collaborative memory field, explain in detail the typical experimental paradigm used in collaborative memory studies, present several mechanistic explanations for collaborative inhibition, and provide an overview and motivation for the studies presented in this dissertation.

1. A History of Group Remembering

The study of group memory has historically been restricted to fields such as sociology, anthropology, and history. Two prominent predecessors of collaborative memory are the collective memory and transactive memory fields. While researchers in these areas study some aspect of group memory, their focus is not typically cognitive in nature. Collective memory is popular within the social science literature and has been associated with the creation of autobiographical memories at cultural and familial levels (Wang, 2008). Typically, collective memory research is approached from a network perspective, rather than at an individual level (Hirst & Manier, 2008). Recently, however, researchers have begun to postulate that collaborative group interactions at smaller scales (such as interactions studied within the collaborative memory field) create the groundwork for the emergence of larger scale collective memories (Choi, Kensinger, & Rajaram, 2017; Maswood & Rajaram, 2019). Conversely, researchers within the transactive memory field are more concerned with the psychology of the individual. Transactive memory is a concept that describes shared memory frameworks between groups of people both large and small (Wegner, 1987). While there is some research within the field that focuses on larger scale community and cultural memory, much of the field is dedicated to empirically studying smaller groups of subjects. Empirical research in this area is typically concerned with shared organizational structure of memory, which is indirectly related to studies of collaborative memory (Wegner, Erber, & Raymond, 1991). While transactive memory studies are based in psychological research, unlike collective memory, the focus is on shared external memory storage within a group more so than the internal cognitive mechanisms underlying group memory.

The current method of studying collaborative memory was adapted from earlier psychological research of individual memory. Originally, collaborative group recall was compared to individual group recall and, unsurprisingly, groups outperformed individuals. Because of this, in the 1950s and 1960s it was believed that group performance followed an additive model. That is, the recall of a collaborative group equals the summed non-redundant recall of the individuals working alone (Lorge & Solomon, 1955). The additive model was supported by several experiments using nonsense materials (Hoppe, 1962; Ryack, 1965), however, was not always representative of results when using meaningful materials. When using meaningful materials, findings suggested that collaborative groups tended to perform at or below pooled individual performance (H. V. Perlmutter, 1953; Stephenson, Abrams, Wagner, & Wade, 1986).

In the early 1990s it was still unclear whether collaborative groups performed better than, equal to, or worse than pooled individuals. Thus, researchers progressed to studying the cross-cuing theory—the idea that group recall might be faciliatory because group members provide each other with helpful cues (Meudell, Hitch, & Kirby, 1992). Unfortunately, researchers were not able to find a faciliatory effect. In fact, Meudell et al. (1992) found an inhibitory effect of group recall and suggested that social factors, such as lowered individual accountability, may be the cause. Andersson and Ronnberg (1995) found that subjects in groups of two may be cross-cuing each other if they were friends but found no such effect if they were strangers. The cross-cuing theory was tested again by B. H. Basden, Basden, and Henry (2000) who once again found no evidence of a faciliatory cross-cuing effect in group recall. With the cross-cuing effect ultimately debunked by numerous studies finding either no significant faciliatory effect or even an inhibitory effect, researchers within the field began to investigate the inhibitory effect of collaboration in more detail.

2. Collaborative Inhibition

The experimental paradigm typically used within the collaborative memory field is an extension of recall paradigms previously used and validated in the field of individual memory. This paradigm involves participants learning a list of words, performing a distractor task individually, and then performing a recall task (typically free recall or cued recall) together in small groups (C. B. Harris, Paterson, & Kemp, 2008). As expected, groups perform better in the recall task than individuals. However, to accurately gauge group performance, collaborative group recall must be compared to nominal group recall (pooled individual recall), not individual recall. In both collaborative and nominal group conditions, subjects learn a list of items individually in the study phase. Then, subjects in collaborative groups are asked to work together with other group members to recall items on the list. The collaborative group response is calculated by counting all unique responses produced by the group. In contrast, subjects in nominal groups are asked to recall items on the list individually and do not recall together. The nominal group response is calculated by counting the total, non-overlapping responses produced by individual group members (see Figure 1.1 for a visualization of this paradigm). When collaborative group recall performance is compared to nominal group recall performance, there is a significant inhibitory effect of collaboration (B. H. Basden et al., 1997; Weldon & Bellinger, 1997)—called *collaborative inhibition*. This is a robust effect that occurs when using a wide variety of study materials, including but not limited to: unrelated word lists (Andersson, Hitch, & Meudell, 2006; Weldon & Bellinger, 1997), categorized word lists (B. H. Basden et al., 2000), story recall (Andersson & Ronnberg, 1995), film clips (Andersson & Ronnberg, 1995; Meudell et al., 1992), and emotional events (Yaron-Antar & Nachson, 2006).

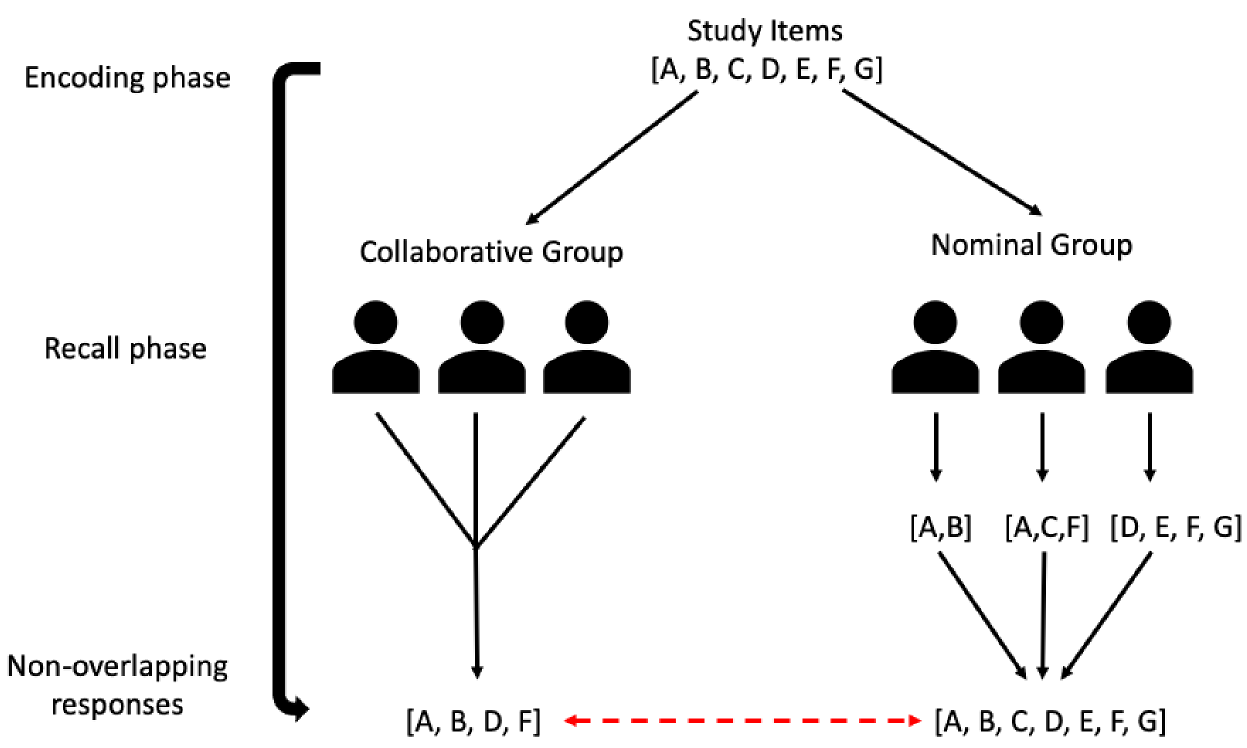

Figure 1.1. *Visualization of collaborative inhibition paradigm.*

3. Mechanistic Hypotheses of Collaborative Inhibition

Currently, there are two groups of mechanistic theories believed to underlie the collaborative inhibition effect: product-based disruption theories and process-based disruption theories. Product-based disruption theories place responsibility on the *product* of recall and point towards items produced by group members interfering with the recall of others. Process-based disruption theories place responsibility on the *process* of recall and point towards group dynamics interfering with recall. Most collaborative memory researchers agree the *product* of recall is more likely to underlie collaborative inhibition (B. H. Basden et al., 1997; Congleton & Rajaram, 2011; Finlay, Hitch, & Meudell, 2000). Several distinct mechanistic theories fall under each category and will be discussed further in this section.

3.1 Product-based disruption theories

Theories that posit the *product* of recall is responsible for collaborative inhibition often originate from theoretical paradigms developed within the individual memory literature. The individual memory analogue to the collaborative inhibition effect is commonly believed to be the part-list cuing effect (Andersson et al., 2006; B. H. Basden et al., 1997; B. H. Basden et al., 2000). Typically, when an individual is asked to use cues to aid recall, their recall performance increases (Tulving, 1974). However, the part-list cuing effect predicts the opposite. When an individual is presented with a random selection of a memorized list as cues, their recall for the remaining words on the list is inhibited (Nickerson, 1984; Slamecka, 1968). Crucially, the part-list cues must be a random subset of the study list for the effect to occur. The similarities between the part-list cuing and collaborative inhibition effects are clear. In both cases, previously studied items are given to a subject as externally produced cues, which then cause disruption in recall. There are three prominent explanations for the part-list cuing effect that may also be responsible for collaborative inhibition: retrieval disruption, retrieval inhibition, and retrieval blocking. All three suggest that the *product* of recall is responsible for the inhibitory effect of collaboration.

3.1.1 Retrieval disruption

The most popular mechanistic hypothesis for collaborative inhibition is the retrieval disruption hypothesis which posits that the inhibitory effects of collaboration occur because individual retrieval strategies are disrupted during group recall (B. H. Basden et al., 1997). This hypothesis was originally used to explain the part-list cuing effect: when randomly chosen, part-list cues interfere with the subject's internal organization of the study list, thus interrupting their idiosyncratic retrieval strategy (D. R. Basden & Basden, 1995). It is theorized that in a collaborative setting, group members provide part-list cues for others in the group—causing collaborative inhibition. According to this hypothesis, each group member develops an idiosyncratic organization of information in memory during the study phase of a recall task which is then disrupted by mismatched cues from other group members when asked to recall in a group. Additionally, once the disruption is removed (i.e., no more group members producing cues), subjects are assumed to remember study items that weren't produced during collaborative recall on subsequent individual recall tasks; that is, memory disruption is not permanent (B. H. Basden et al., 2000; Congleton & Rajaram, 2011).

B. H. Basden et al. (1997) were the first to provide experimental evidence supporting the retrieval disruption hypothesis as an explanation for collaborative inhibition. Under this hypothesis, they predicted that the inhibitory effect of collaboration would be reduced when study materials were more organized, due to less room for idiosyncratic organization within subcategories. To test this, B. H. Basden et al. (1997) compared collaborative recall performance on two organized word-lists: one list consisted of 15, six-word categories and the other list consisted of 6, fifteen-word categories. As predicted, when individuals learned items from an organized list with 15, six-word categories, collaborative inhibition was reduced because there was enough overlapping structure within the study materials to prevent much idiosyncratic subcategory organization. Further work by B. H. Basden et al. (2000) suggests that after recalling once in a collaborative group, administering a second individual recall caused group members to recall more words. If recall is inhibited during group recall but released during subsequent individual recall, then the inhibition is temporary—as would be the case if retrieval disruption were responsible for collaborative inhibition. This effect is also seen in the part-list cuing literature when the part-list cues are removed in a subsequent recall trial and inhibition is released (D. R. Basden & Basden, 1995).

Work by Finlay et al. (2000) expanded upon the work by B. H. Basden et al. (1997) and provided further evidence supporting retrieval disruption as a driving factor for collaborative inhibition. The predictions they tested were (1) that collaborative inhibition is temporary, (2) collaborative inhibition should be reduced after collaborative encoding, and (3) collaborative inhibition should be reduced or non-existent in cued recall. If retrieval disruption were the driving mechanism causing collaborative inhibition, then (1) inhibition should only occur during collaborative recall and performance would return to normal in post-collaborative individual recall, (2) collaborative encoding would reduce the differences in idiosyncratic memory organization, thus decreasing inhibition during recall, and (3) cued recall would prevent group members from providing incongruous cues which disrupt individual search strategies. They found that all three predictions were supported by their experimental data, thus providing further evidence in favor of the retrieval disruption hypothesis.

3.1.2 Retrieval inhibition

A second possible mechanistic explanation for collaborative inhibition is retrieval inhibition which posits that strengthening of cue words inhibits the memory for non-cued words by suppressing memory representations, which prevents those words from being retrieved (Bäuml & Aslan, 2004). In a collaborative setting, words that are cued by group members would be strengthened in memory and words that are not recalled by the group would be weakened, causing permanent suppression of unrecalled words for all group members. It is also important to note that this memory impairment should persist after collaboration regardless of the method in which memory is cued. That is, the impairment should also be noticeable in post-collaborative free recall and recognition tests (Bäuml & Aslan, 2006).

The body of experimental evidence for this hypothesis is relatively smaller than that of the retrieval disruption hypothesis. However, several recent studies have found supporting evidence for retrieval inhibition, by observing an incomplete release form inhibition during post-collaborative individual recall tasks (which is not predicted by retrieval disruption). On post-collaborative individual tasks, subjects often forget to recall words they contributed to earlier collaborative recall (Blumen & Rajaram, 2008) and this effect increases as the effect size of collaborative inhibition increases (Congleton & Rajaram, 2011), suggesting a long-lasting detrimental effect of collaboration.

It has historically been difficult to examine post-collaborative effects due to the competing nature of retrieval inhibition and beneficial effects of re-exposure during collaboration. It is possible that retrieval inhibition and post-collaborative forgetting occur but are offset by re-exposure effects during collaborative recall. Barber, Harris, and Rajaram (2015) avoid this entanglement by having participants study non-overlapping lists, which prevents re-exposure benefits, before performing collaborative recall. They found that without re-exposure benefits, recall remained inhibited on subsequent individual free recall and recognition tests. However, the post-collaborative inhibition effect was reduced compared to collaborative inhibition. These findings suggest that both retrieval inhibition and retrieval disruption may be responsible for collaborative inhibition.

3.1.3 Retrieval blocking

A third possible mechanism for collaborative inhibition is retrieval blocking which posits that cue words become stronger candidates for retrieval and people are more likely to recall previously cued words as opposed to new words. During the part-list cuing task, participants are more likely to think of the cued words which blocks access to non-cued words (Rundus, 1973). In a collaborative setting, group members would be more likely to recall previously recalled words than produce novel responses. While retrieval blocking prevents access to non-cued words, the memory representation itself is not suppressed. That is, the memory deficit should remain when memory search is self-guided but be eliminated when cues are provided, such as on a recognition test. Barber et al. (2015) found that inhibition was present in both post-collaborative free recall and recognition tests, suggesting that retrieval blocking may not play an active role in collaborative inhibition.

3.2 Process-based disruption theories

Theories that suggest the *process* of recall is responsible for collaborative inhibition often originate from the social psychology literature. These theories posit that there's something about the collaborative process itself (independent of the items produced) that inhibits group performance. While social group processes (such as brainstorming) are certainly similar in some ways to collaborative recall, these theories have almost entirely been debunked as mechanistic explanations for collaborative inhibition by cognitive researchers.

3.2.1 Social factors

In any context where group interaction occurs, it seems intuitive that social interaction and motivation would be important. Therefore, it follows that collaborative inhibition could be caused by a variety of social factors. One such social factor is social loafing, the tendency for group members to not work as hard in a group setting as they would have alone (Latane, Williams, & Harkins, 1979). Social loafing as a possible mechanistic explanation for collaborative inhibition is implied by previous group research in a wide variety of fields that show a similar loss of individual productivity. These areas include bystander intervention (Latane & Nida, 1981), physical activities such as rope-pulling (Ingham, Levinger, Graves, & Peckham, 1974), and cognitive tasks like brainstorming (Diehl & Stroebe, 1987; Taylor, Berry, & Block, 1958). Given the similarity between brainstorming and collaborative recall tasks, it is possible that the same mechanisms could be responsible for both inhibitory effects. However, while there is some evidence that social loafing may play a detrimental role in brainstorming activities (Diehl & Stroebe, 1987), the experimental evidence available is not enough to account for the whole effect.

Social and motivational factors have been unable to account for the inhibitory effect within collaborative memory. Experiments that manipulated factors important for social loafing, such as personal accountability and individual motivation, were not able to decrease collaborative inhibition (Weldon, Blair, & Huebsch, 2000). However, collaborative inhibition *can* be decreased or even eliminated by manipulating cuing or the order in which items are learned (Andersson et al., 2006). If social factors were the sole cause of collaborative inhibition, such manipulations would not decrease the inhibitory effect. The inability to reduce collaborative inhibition by eliminating possible social factors suggests that other mechanisms may be responsible for the effect.

3.2.2 Production blocking

The production blocking hypothesis posits that waiting to contribute while other group members produce responses blocks the ability to recall information (Diehl & Stroebe, 1987). For example, while recalling in a group, individuals might forget their response while waiting for other group members to finish talking. Thus, the cause of collaborative inhibition would not be because of the responses produced by the group but because the process of collaboration forces participants to wait to produce responses. The loss of productivity during brainstorming has been documented since the mid-20th century (Bouchard & Hare, 1970; Taylor et al., 1958). Diehl and Stroebe (1987) ran a series of experiments investigating the possible cause of productivity loss within brainstorming groups and tested three separate hypotheses: evaluation apprehension, free riding, and production blocking. They found very little evidence in support of evaluation apprehension or free riding, both of which are social factors, and concluded that production blocking, the cognitive explanation, was the most likely. However, most studies involving production blocking during collaborative and individual recall have either concluded that there isn't enough supporting evidence to fully account for collaborative inhibition or that production blocking can't be the cause of collaborative inhibition (Andersson et al., 2006; Finlay et al., 2000; Wright & Klumpp, 2004).

In an individual part-list cuing task, Andersson et al. (2006) found that the inhibitory effect of cues during recall extended beyond cue presentation, indicating they do more than disrupt retrieval operations at the time of presentation. This pattern is evidence against a simple blocking effect since the cues inhibited recall even after they had been presented. To further tease apart the cause of recall inhibition in a part-list cuing task, Andersson et al. (2006) compared part-list cues with extra-list cues when they were presented at the beginning of recall or throughout recall. While both cue sets are equivalent when it comes to blocking power, they may be different when it comes to retrieval disruption abilities. Theoretically, the part-list cues are more likely to disrupt retrieval strategies because they are associated with the study list items whereas the extra-list cues are not. The authors found that both cue types caused inhibition if presented at the beginning of recall, but part-list cues caused significantly more inhibition when the cues were presented throughout recall. These results suggest that there is not enough evidence to attribute the part-list cuing effect entirely to production blocking.

While production blocking wasn't ruled out as a possible explanation behind the part-list cuing effect, Wright and Klumpp (2004) demonstrated that production blocking is most likely not responsible for collaborative inhibition. They compared nominal group performance to two conditions of collaborative recall: the seeing condition and the not-seeing condition. In both conditions, collaborative groups consisted of two members and were instructed to use the turn-taking method of recall. Group members in the seeing condition passed the same sheet of paper back and forth and wrote down responses on their turn; they were able to review responses produced by their partner before writing their own response. Group members in the not-seeing condition were unable to see responses produced by their partner before writing their own response. To achieve this, both group members had their own, separate pieces of paper that their partner never saw. Group member 1 recalled an item and wrote it down on their piece of paper. Then, group member 2 was allowed to respond and wrote a response down on their own piece of paper. Additionally, there was a divider between group member 1 and 2 during the entire collaborative recall task. Thus, group members could not see what other responses were produced but still had to wait for their partner to produce their responses. If production blocking were responsible for collaborative inhibition, participants assigned to the second collaborative condition should still experience inhibition. However, this was not the case, suggesting that production blocking is not responsible for collaborative inhibition.

3.3 Theoretical relevance to current work

As discussed in the previous section, there are many possible mechanisms with varying amounts of experimental support, however, mechanisms suggesting the *product* of recall interferes with output are generally accepted over process-based mechanisms. One goal of this dissertation is to develop and employ a collaborative modeling framework to further investigate possible mechanisms responsible for collaborative inhibition. Given the nature of this collaborative framework and the findings from the experimental literature, the focus of this dissertation will be on various product-based theories of disruption such as retrieval disruption, retrieval inhibition, and retrieval blocking.

4. Experimental Factors Affecting Collaborative Inhibition

As of now, collaborative inhibition is the predominant focus of research within the field of collaborative memory. This is a robust effect that occurs when using a wide variety of study materials. However, there is also experimental evidence that other factors, beyond study materials, moderate the severity of the collaborative inhibition effect including: category size, encoding conditions, memory task type, memory type (episodic vs. semantic), age, relationships between group members, expertise, and group size. These factors will be discussed in more detail in the following sections.

4.1 Category size

The motivation for testing category size differences comes from similar tests within the part-list cuing literature. D. R. Basden and Draper (1973) compared the effect size of part-list cuing on individual recall from categorized lists. The first list had 15 categories with 6 words each and the second list had 6 categories with 15 words each. They found that individual recall was worse when receiving part-list cues from the list with fewer categories. They argued that individuals were more likely to create within-category organization strategies for the larger categories, thus, the part-list cues were more likely to interfere with each individual's retrieval strategy.

An experiment very similar to the above was performed within the collaborative memory literature. B. H. Basden et al. (1997) compared collaborative and nominal recall on categorized lists: one list with 6 categories that had 15 words in each and another list with 15 categories that had 6 words each. Their results were akin to those in the analogous part-list cuing experiment. They found that the performance of collaborative groups varied by list type. That is, collaborative groups performed worse when category sizes were larger. Just like the results from the part-list cuing experiment, B. H. Basden et al. (1997) suggest the reason collaborative groups perform worse when categories were bigger is because there is a higher likelihood for subcategory organization which is unique to each group member. When group members then provide cues for each other during recall, the idiosyncratic retrieval strategy of each individual is disrupted. Thus, these results provide support for the retrieval disruption hypothesis.

4.2 Collaborative encoding

There are mixed results when it comes to the effect of collaborative encoding on subsequent collaborative and individual retrieval. According to the retrieval disruption hypothesis, if collaborative encoding were to align group members' internal organization of study materials, then collaborative inhibition should decrease during collaborative retrieval. B. H. Basden et al. (1997) found that similar strategies for encoding tend to produce similar orders of output during retrieval. If the retrieval output order is similar enough, the retrieval disruption from incongruous cues is reduced, effectively eliminating collaborative inhibition. Finlay et al. (2000) designed a collaborative encoding task such that group members organization of study materials were more similar by ensuring participants learned the study materials in the same order. They found that collaborative inhibition was reduced in this case, thus supporting the retrieval disruption hypothesis.

On the other hand, there have also been studies that show the opposite effect and suggest that there may be a collaborative encoding deficit that is independent of collaborative inhibition. Andersson and Ronnberg (1995) found that collaborative encoding not only inhibited subsequent collaborative recall but also negatively affected subsequent individual recall. In this case, it was believed that collaborative encoding caused subjects to use cues generated by their group members when they were asked to recall the study items. Earlier research suggests that self-generated cues are better retrieval cues than cues generated by others (Mantyla & Nilsson, 1983), thus, using cues generated by group members would negatively affect subsequent recall. Additionally, Barber, Rajaram, and Aron (2010) tested the effect of collaborative encoding on a subsequent cued recall test. Collaborative inhibition has been shown to disappear during cued recall tests (Finlay et al., 2000), thus these recall conditions were specifically tailored to tease apart a collaborative encoding deficit from the collaborative inhibition effect. Results suggested that the collaborative encoding deficit is separate from the collaborative inhibition effect as collaborative encoding had negative effects on individual, nominal, and collaborative recall despite employing a cued recall test.

4.3 Memory task type

Collaborative inhibition is a robust effect when it comes to episodic free recall of many different study materials (Andersson et al., 2006; Andersson & Ronnberg, 1996; B. H. Basden et al., 1997), however there is conflicting evidence concerning the presence of collaborative inhibition in other memory task types. Early research investigating collaborative inhibition in memory tasks such as recognition memory, cued recall, or semantic memory do not show collaborative inhibition (Andersson & Ronnberg, 1996; Finlay et al., 2000). Clark, Hori, Putnam, and Martin (2000) found that collaboration in recognition memory tasks results in a faciliatory effect when compared to nominal groups. More specifically, they found that there was a consistent faciliatory effect for recognizing targets but not for rejecting distractors. In a similar pattern, early studies investigating the effect of collaboration in cued recall have shown that there is no significant faciliatory or inhibitory effect (Finlay et al., 2000). It was hypothesized that the discrepancy between the memory task types is because external cues are provided for the subjects during cued and recognition memory which do not interfere with their internal cue production like they do in free recall (Rajaram & Pereira-Pasarin, 2010).

Recently, collaborative inhibition has been shown in both recognition tasks (Danielsson, Dahlström, & Andersson, 2011) and cued recall tasks (Kelley, Reysen, Ahlstrand, & Pentz, 2012; M. L. Meade & Roediger, 2009). It is unclear why there is a discrepancy between the earlier research and more recent research, however, it could point towards mechanistic processes other than retrieval disruption at play during collaborative recall.

4.4 Episodic vs. semantic memory

In addition to memory task type, there also seems to be differences in how collaboration affects semantic memory and episodic memory recall tasks. Typically, collaborative inhibition is seen when subjects are asked to recall episodic memories. When Andersson and Ronnberg (1996) tested groups of friends and non-friends on both episodic and semantic retrieval tasks, they found that inhibition was dependent on the memory type. Only explicit, episodic tasks were negatively impacted by collaboration while implicit semantic tasks were not. The authors conclude that this result is logical because semantic memory is typically more organized than episodic memory and does not need to be cued with as much precision to retrieve information.

The task used to prompt semantic recall is arguably not a good semantic equivalent to episodic free recall. Andersson and Ronnberg (1996) asked subjects to recall generic Swedish history questions to test semantic recall. While history questions certainly probe semantic memory, it is unclear how much knowledge is shared between subjects and what the internal organizational structure would look like for history facts. A much more suitable task for testing collaborative inhibition in semantic memory would be a verbal or letter fluency task. The semantic knowledge probed in fluency tasks is relatively simpler, is more likely to be shared across all subjects, and subjects have been shown to produce responses in clusters—indicating internal organization of semantic memory plays a large role in these tasks (Troyer, Moscovitch, & Winocur, 1997). **Chapter 4** presents the results of a behavioral experiment comparing performance of collaborative and nominal groups on a verbal fluency task. Additionally, the verbal fluency data is analyzed using optimal foraging models for differences in search methods between groups.

4.5 Age

Few collaborative memory studies have included age as an experimental variable. However, studies that have focused on age have shown that collaborative inhibition is present in all age groups. Andersson (2001) predicted that collaborative groups consisting of younger individuals would be less able to prevent collaborative inhibition than groups of older individuals. This prediction is based on developmental stages related to perspective taking. The younger the child, the more difficult perspective taking is (Piaget, 1959) which could feasibly make collaborative inhibition worse. To test this prediction Andersson (2001) compared group recall between pairs of 7-year-olds and pairs of 15-year-olds. They found that their prediction was correct and the 7-year-old pairs were more affected by collaborative inhibition that the 15-year-olds. Though the two groups showed different levels of inhibition, the effect did not disappear in the 15-year-old pairs.

Additionally, there have been a few studies investigating the effect of collaboration on older adults. Johansson, Andersson, and Ronnberg (2000) tested elderly couples, pairs of elderly strangers, and elderly nominal groups on retrospective and prospective memory tasks. They found that both the elderly couples and pairs of elderly strangers performed worse on both tasks than the elderly nominal groups. In another study, Henkel and Rajaram (2011) directly compared the collaborative recall performance of young adults and older adults on a categorized word list. They found that while both groups exhibited collaborative inhibition, the groups of older adults consistently underperformed compared to the younger adults. This outcome is not surprising given that older adults have previously been found to perform worse in free recall tasks compared to younger adults (Hultsch, 1971; M. Perlmutter, 1979). The results of the few studies focusing on age-related factors in collaborative memory have shown that while younger children may be more adversely affected by collaboration than older children or adults, collaborative inhibition is a persistent effect across all age groups.

4.6 Shared background knowledge

4.6.1 Relationships

Proponents of transactive memory suggest that people in close relationships develop a shared, transactive memory which is greater than either person's individual memory. Wegner et al. (1991) found that couples who had been together for at least 3 months were able to outperform stranger-pairs on unstructured memory tasks. However, when subjects were told which parts of the task each person should remember (instead of letting the division of remembering occur naturally), couples underperformed compared to stranger-pairs. These findings from the transactive memory literature are the basis for relationship research within the collaborative memory field. Based on the outcome of this experiment, if subjects have a pre-existing relationship, they may be able to reduce or eliminate collaborative inhibition.

Within the collaborative memory literature, there have been a few studies comparing friend-pairs with stranger-pairs (Andersson & Ronnberg, 1995, 1996) and married couples with stranger-pairs (Johansson et al., 2000) that show a reduction in collaborative inhibition. However, while there was some reduction in collaborative inhibition, the effect was not completely eliminated. The slight advantage of friends and married couples compared to stranger-pairs may be explained by social factors or cognitive factors. It is possible that people with pre-existing relationships have more practice communicating and interacting with each other which may reduce collaborative inhibition. It may also be the case that people with pre-existing relationships have more shared background knowledge and experience which could result in a more similar organization of study list items. It is difficult to tease apart these two explanations in a behavioral experiment but is possible to study more closely in a collaborative model.

4.6.2 Expertise

Similarly, it is unclear whether shared background knowledge or social factors are responsible for collaborative facilitation in collaborating experts. M. L. Meade, Nokes, and Morrow (2009) found that if subjects had expertise in the area of study, collaboration was beneficial compared to nominal groups. This effect was discovered when they compared the performance of expert pilots, novice pilots, and non-pilots on a collaborative memory task involving aviation scenarios. They found that while the non-experts still suffered from collaborative inhibition, the expert pilots benefited from collaboration during recall.

This finding is supported by the retrieval disruption hypothesis. If encoding strategies are similar enough, then retrieval outputs tend to be more similar. In this case, retrieval disruption from incongruous cues is reduced, effectively eliminating collaborative inhibition (B. H. Basden et al., 1997; Finlay et al., 2000). Experts encode information differently than novices (Herzmann & Curran, 2011; Morrow et al., 2008), and the authors believe it is likely that the pilots encoded the study material in a similar manner, thus minimizing collaborative inhibition. As for the faciliatory effect of collaboration, it remains unclear if this effect is unique to certain types of expertise. M. L. Meade et al. (2009) mention that aviation expertise includes in depth training in communication skills which may be responsible for the faciliatory effect they found. It is uncertain whether collaborative facilitation in experts is unique to certain domains of expertise or is generalizable to all experts. Unfortunately, there is no available research comparing different forms of expertise to address this question.

Both previous relationships and shared expertise mitigate collaborative inhibition to some degree. However, in both cases, it is unclear whether social factors, such as effective communication, or cognitive factors, such as shared background knowledge are responsible. A collaborative model would be able to tease apart the effect of shared background knowledge and social factors. That is, a collaborative framework would not include social factors that can interfere or compound with the effect of shared memory organization in separate models. **Chapter 3** presents a modeling experiment exploring just this topic: the effect of shared background knowledge on collaborative inhibition without interference from social factors.

4.7 Group size

Several studies suggest that as group size increases, the magnitude of the collaborative inhibition effect increases. Within the brainstorming literature, productivity loss from collaboration does not occur until the group size exceeds two (Diehl & Stroebe, 1987). In fact, there is evidence that collaborative and nominal brainstorming groups of size two do not differ significantly in idea output (Paulus & Dzindolet, 1993). However, when Bouchard and Hare

(1970) tested the idea output of five, seven, and nine person brainstorming groups, they found the magnitude of the productivity loss increased as group size increased.

Within the collaborative memory literature, there is some conflicting evidence concerning collaborative inhibition in groups of size two. Usually, collaborative inhibition is found in pairs (Andersson & Ronnberg, 1996; Finlay et al., 2000), but there are some studies which show no significant inhibitory effect of collaboration in such small groups (B. H. Basden et al., 2000; Meudell, Hitch, & Boyle, 1995). On the other hand, collaborative inhibition has consistently been shown to effect groups of three (B. H. Basden et al., 1997; Blumen & Rajaram, 2008; Weldon & Bellinger, 1997). Additionally, studies that compare groups of three to groups of four have found that larger groups are relatively more affected by collaborative inhibition (B. H. Basden et al., 2000). These results are predicted by the retrieval disruption hypothesis—as more group members are added, there is a greater chance for each individual to have their retrieval strategy disrupted because more group members are providing external cues.

Contrary to previous findings, Gates, Suchow, and Griffiths (2022) found no significant interaction between recall method (nominal vs. collaborative) and group size (2, 3, 4, 8, 16) in either turn-taking or free-for-all recall methods. This is the first experimental study that tested group sizes larger than 4, however these findings are contradictory to both previous experimental research and predictions of the retrieval disruption hypothesis. The lack of collaborative inhibition in this study is surprising and encourages a closer look at methodological differences in this study compared to other collaborative memory studies. First, unlike most other behavioral studies, subjects were recruited online via Amazon Mechanical Turk and interacted virtually via a chatroom. Subjects were able to read other's responses as they were typed but there was no permanent list for subjects to reference during recall. It is unclear how much attention subjects gave to group member’s responses during recall. If subjects did not use group members’ cues, then collaborative inhibition would disappear. Second, subjects were not prevented from repeating the online study. The authors note that 30% of their data for the free-for-all experiment was generated by participants that repeated the task. They found that while there was no improvement for participants repeating the nominal task, participants repeating in the collaborative condition did show improvement. Thus, subjects repeating and improving on the collaborative task may also contribute to the lack of collaborative inhibition. While there are several factors of this study design that could have mitigated collaborative inhibition, the results of this study undeniably raise questions about the predictions the retrieval disruption hypothesis makes and emphasizes the need for more studies explicitly investigating group size effects.

Group size is another manipulation that lends itself to analysis with a collaborative memory model, as increasing group size is much easier in a modeling experiment than a behavioral experiment. **Chapter 3** presents a modeling experiment investigating the effect of increasing group size on a collaborative inhibition.

5. Overview of Dissertation

As of now, most research in the collaborative memory field is behavioral and focuses on episodic memory tasks. However, there are limitations to behavioral research that can and should be supplemented by modeling efforts. While there is an abundance of behavioral support for the retrieval disruption hypothesis and several experimental factors that modulate the effect, there is growing evidence of a multi-process account of collaborative inhibition in which other mechanisms may be involved, such as retrieval inhibition, retrieval blocking, and memory homogenization. There are currently no cognitive modeling efforts focused on teasing apart the contribution of each mechanism to collaborative inhibition and a surprising lack of studies investigating the effect of collaboration on semantic memory tasks. Additionally, memory homogenization, the idea that group member's memories become more similar during retrieval, is another possible cause of collaborative inhibition that can only be explicitly studied using models of collaborative memory, as directly measuring memory similarity in behavioral experiments is impossible. The studies presented in this dissertation were completed for the purpose of expanding the modeling efforts within the field of collaborative memory for both episodic and semantic memory tasks.

Chapter 2 describes the motivation and implementation of a collaborative modeling framework for studying collaborative inhibition in episodic free recall. Analyzing how a collaborative model produces collaborative inhibition will provide valuable insights into which mechanisms may be contributing to the effect. **Chapter 3** investigates extensions of the collaborative modeling framework with the goal of using the model to understand cognitive mechanisms responsible for collaborative inhibition and testing hypotheses which would be difficult or nearly impossible to investigate via behavioral methods. Such extensions include investigating the effect of memory homogenization, shared background knowledge, group size, and post-collaborative effects on collaborative inhibition. **Chapter 4** analyzes collaborative verbal fluency data collected for this dissertation using optimal foraging models of semantic memory previously shown to be useful for understanding memory search behavior in the verbal fluency task (Hills, Jones, & Todd, 2012). Finally, **Chapter 5** summarizes insights gained in this dissertation and presents a discussion of several future applications for modeling collaborative inhibition.

CHAPTER 2

THE COLLABORATIVE SAM (cSAM) FRAMEWORK

Research focusing on social media information presently dominates the group behavior literature, and includes topics such as community identification, “fake news” detection, topic modeling, and misinformation prevention. This research stems from the fields of network science and linguistics and tends not to incorporate or consider cognitive mechanisms in their models. Until now, the only attempt at modeling collaborative memory was made by Luhmann and Rajaram (2015) whose main goal was to model information transmission at network-scale by taking an agent-based modeling approach. Though their main goal was not to model collaborative inhibition, during the verification phase of their model, the authors were able to produce collaborative inhibition when groups of three agents were tasked with performing collaborative recall. Additionally, they were able to model the effect of group size on collaborative memory. However, while this model included psychologically based agents that were able to encode and retrieve memories, the main goal of the study was to examine the effect of information transmission on network behaviors. This model did not aim to synthesize individual and collaborative memory processes, the broader goal of the current chapter.

1. Modeling Collaborative Memory

While a large body of experimental research and verbal conceptual frameworks exist for collaborative memory, there are currently no formal computational models to guide the field. The present research adapts a prominent model of individual memory recall, the Search of Associative Memory model (SAM; Raaijmakers & Shiffrin, 1980; Raaijmakers & Shiffrin, 1981), to predict the results of collaborative recall. Most relevant, SAM has been used to explain the mechanisms responsible for part list-cuing (Raaijmakers & Shiffrin, 1981). The processes explaining part-list cuing and collaborative inhibition might be similar: part-list cuing shows a lowering of recall when a subset of randomly selected list items are provided by the experimenter as cues – the random cues harm recall because they do not match the subjective organization in memory. In collaborative recall, the lowering of group performance could be due to the use of recalls provided by others in the group as cues: a cue provided by another group member is effectively 'random' in that it is unlikely to match the subjective organization of the group member who is induced to use it. SAM is used to model individuals who free recall word lists independently; it uses the standard SAM assumption that the most recent word recalled is used as the next cue for searching memory. A reasonable and slight adaptation of SAM was used, termed cSAM, to model group recall; cSAM assumes that every member of the group uses the most recent word recalled by any group member as the next cue for searching memory. With this exception, the same model is used with the same parameter values. Our simulations support the retrieval disruption hypothesis: that the cues produced by the group tend to mismatch the subjective organizations of the group members. The results provide a formal framework unifying individual and collaborative memory research.

The motivation for using SAM over other possible cognitive models is as follows. First, SAM is well-studied and is the most widely used model in episodic memory research (Wilson &

Criss, 2017; Wilson, Kellen, & Criss, 2020). Second, SAM is one of the only cognitive models that has been shown to successfully model the part-list cuing effect in individual memory (Raaijmakers & Shiffrin, 1981). Finally, the architecture of the model affords a coherent framework to extend to multiple models working collaboratively. If SAM can be modified to explain collaborative phenomena without changing the fundamental architecture, any of the SAM models in isolation would still retain the explanatory power for the range of behavioral phenomena in individual memory paradigms thus producing a unified account of both individual and collaborative memory phenomena.

1.1 Search of associative memory (SAM)

SAM is a cue-dependent probabilistic search theory of retrieval and is typically applied to simulations of free and cued recall. The model makes use of a two-stage memory system: short-term memory and long-term memory. The short-term memory system is where processes such as encoding and rehearsal are carried out during study and where retrieval is controlled during testing. Long-term memory contains traces represented as an association matrix of study items and environmental context (context-item information) and item to item-plus-context information (item-item information). Context information represents information available during encoding that identifies the context of the list rather than any specific item, such as emotions, sensations, or environmental details.

The traces in long-term memory are formed and stored during the time that items are present in short-term memory in a limited capacity rehearsal buffer. The items present in the buffer together at each moment are determined probabilistically, so that different participants form different subjective organizations (different associative strengths) in long-term memory. The associative strength between a context cue and a context-item trace is a linear function of the time that item was rehearsed in short-term memory. The associative strength between an item-plus-context cue and an item-item trace that contains the same item (i.e. item-to-self relationship) is a different linear function of the time that item was rehearsed in short-term memory. The associative strength between an item-plus-context cue and an item-item trace containing a different item (i.e. item-to-other relationship) is a third linear function of the time the two items were together in the rehearsal buffer. In many tasks, including the present ones, the contents of short-term memory are cleared by a distracting task such as arithmetic before recall begins.

Retrieval from long-term memory is then carried out by probing memory either with a context cue alone (at the start of retrieval or when an item cue is no longer helpful) or with an item-plus-context cue. Learning during retrieval is represented by additions to the associative strengths between cues and traces when a successful recall occurs. Recall due to context only cuing increments strengths from the context cue to that trace and the strength of the item to its own trace. Recall due to item-plus-context cuing increments those strengths and increments the strength of the item cue to the trace containing the recalled item. Table 2.1 gives a brief description of the standard parameters included in SAM. The first 11 parameters are for SAM applied to both nominal and collaborative groups free recalling uncategorized lists. The last two parameters apply to models free recalling from categorized lists. The columns to the right provide the parameter values used to predict collaborative inhibition for nominal and collaborative groups, for both uncategorized and categorized lists (as explained in detail later).

Table 2.1. *SAM Parameter Descriptions*

Parameter	Description	Uncategorized Values	Categorized Values
t	Presentation time per word during encoding	2s	2s
r	STM buffer size	4	4
a	Weight for context to word association during encoding	.08	.07
b	Weight for word to other word association during encoding	.08	.07
c	Weight for word cue to same word association during encoding	.08	.07
d	Associative strength for words not appearing in buffer together	.02	.02
e	Incrementing parameter for context-to-word association during retrieval	.7	.7
f	Incrementing parameter for word-to-word association during retrieval	.7	.7
g	Incrementing parameter for word-to-self association during retrieval	.7	.7
Kmax	Number of retrieval failures that end retrieval	30	30
Lmax	Number of retrieval failures before returning to context cues	3	3
h	Starting association for words in the same category (categorized)	--	.25
i	Starting association for words in different categories (categorized)	--	.005

1.1.1 Retrieval in SAM

As mentioned, short-term memory is cleared in the present tasks, so that all recall comes from long-term memory. Retrieval begins by probing memory with a context cue. Traces are activated in proportion to their associative strength to the context cue, and one trace is sampled in proportion to that strength. The probability of recalling the item in the sampled trace (termed recovery) rises with the associative strength. If recall fails, either because a word is not recovered from the sampled trace, or because a word is recovered but had been recalled previously, another sample is made and this continues until an item is recalled, or until *Kmax* total accumulated failures occur (*Kmax* is never reached at the start of recall, but eventually is reached and terminates recall).

If an item is recalled, then that item plus context is used next to probe long-term memory. Traces are activated in accord with the associative strength between the item and context probe and each trace. A trace is sampled in proportion to its strength. The probability of recalling the item in that trace rises with the associative strength between the cue and trace, but successful recall only occurs if that item had not already been recalled (which counts as a failure). If recall fails another sample is made and this continues until a new item is recalled, in which case the new item plus context is used to probe memory, or until *Lmax* is reached. If *Lmax* is reached, then retrieval ends if *Kmax* failures have accumulated overall; if *Kmax* has not been reached, then the retrieval cue is changed to a context only probe. In this way, the memory search continues until *Kmax* total failures accumulate, at which point recall stops.[1]

Equation 2.1a gives the probability of sampling a context-word trace, W_{iS}, using only context, C_T, as a memory probe. Equation 2.1b gives the probability of sampling a context-word trace, W_{iS}, given both context, C_T, and a word cue, W_{kT}, as a memory probe. The S subscript indicates the item as it is stored in memory.

$$P_S(W_{iS}|C_T) = \frac{S(C_T, W_{iS})}{\sum_{j=1}^{n} S(C_T, W_{jS})} \tag{2.1a}$$

$$P_S(W_{iS}|C_T, W_{kT}) = \frac{S(C_T, W_{iS})S(W_{kT}, W_{iS})}{\sum_{j=1}^{n} S(C_T, W_{jS})S(W_{kT}, W_{jS})} \tag{2.1b}$$

[1] If *Kmax* is reached while search is using a word cue and *Lmax* has not yet been reached, then search continues until a new word is recalled and used, or until *Lmax* is reached and then search ends. This happens very seldom, not enough to alter any of the present simulation results significantly. Also, the original version of SAM did not end search when *Kmax* was reached but had a process of 'rechecking', using words recalled earlier as cues. Rechecking was not used in any of the present modeling. As cSAM is described it will become clear that rechecking added to both the nominal and collaborative groups would increase the degree of collaborative inhibition predicted by inducing the use of ineffective cues by the collaborative group. Rechecking was not included because the basic SAM predicts collaborative inhibition without it.

Once an item is sampled from memory, the recovery process begins. Equation 2.2a shows the probability of recovering the word, W_i, in the sampled trace given only context as a memory probe. Equation 2.2b shows the probability of recovering the word in the sampled trace given both context and a word cue as a memory probe. The recovered word is counted a failure if already recalled and otherwise a success.

$$P_R(W_i|C_T) = 1 - exp\{-S(C_T, W_{iS})\} \tag{2.2a}$$

$$P_R(W_i|C_T, W_{kT}) = 1 - exp\begin{Bmatrix} -S(C_T, W_{iS}) \\ -S(W_{kT}, W_{iS}) \end{Bmatrix} \tag{2.2b}$$

1.1.2 Part-list cuing in SAM

In the original formulation of the SAM model, Raaijmakers and Shiffrin (1981) were able to successfully model the part-list cuing effect. Typically, the SAM model uses internal cues to perform free recall. In the case of part-list cuing, the model must first use external cues, provided by the experimenter, to perform recall before transitioning to internal cues once the external cues have been used. The first simulation compared recall performance of a control model, which was given no external cues, and a cued model, which was given external cues. Both models were trained on lists of size 30 and with a presentation time of 2 seconds per item. The cued model was given 15 cues during recall. Additionally, the interitem strength parameter of the models was varied during this experiment to better adhere to the original experiment detailing the part-list cuing effect (Slamecka, 1968). When the control model and the cued model performances were compared, Raaijmakers and Shiffrin (1981) found that the cued models recalled fewer critical items (non-cue items) across all values of the interitem strength parameter, as was predicted by the part-list cuing effect.

This result is important not only because the cued group performed worse than the control group but because the performance hit occurred despite a factor that aids the cued group. A recovery rule of the SAM model makes it such that the probability of recovery after a memory item is sampled using an item-plus-context cue is greater than when a memory item is sampled using a context only cue. That is, during the memory search process, the model can either use context cues or previously recalled item cues to further aid in memory search. In the original model implementation, the model would typically begin free recall by using a context cue (stored during the study phase) to probe its long-term memory storage. Then, once the model successfully recovers an item from memory, that item is then used in conjunction with context as the next memory probe. So, it is somewhat surprising, given the fact that the recovery rules of the model favor item-plus-context cues, that the cued models would consistently perform worse than the control models.

Raaijmakers and Shiffrin (1981) had a few possible explanations for the presence of the part-list cuing effect in SAM. First, they suggest the item-to-context strength may be playing a role. Both the control and cued models use cues, but the control model uses only self-generated cues while the cued model uses cues provided by the experimenter first. Since both models extensively use interitem associative cuing, the overall performance of the models should be comparable. Because of the different types of cuing, self-generated and experimentally provided, it follows that in the case where context alone is unable to produce useful cues, the control model should perform worse than the cued model. When Raaijmakers and Shiffrin (1981) compared model performances with a low value of the item-to-context parameter, which controls the associative strength between memory items and context, they found that the control model underperformed compared to the cued model. And, as predicted, when the value of the item-to-context parameter was increased the typical part-list cuing effect reappeared in the models.

A second explanation for the effect has to do with the selection of the part-list cues. Specifically, the part-list cues must be a random sample of the items from the study list. If they are not, then the cues can be beneficial to the cued model. This pattern has also been found in experimental data. Typically, when a subject is asked to use cues to aid recall, they perform much better than without the cues (Tulving, 1974) as the cues are able to increase recall for highly associated words (Roediger, 1978). Raaijmakers and Shiffrin (1981) suggest that if the cues presented during a part-list cuing task were aligned with a subject's idiosyncratic subcategorization of the study materials then the part-list cues could greatly benefit recall. This prediction was later confirmed by D. R. Basden and Basden (1995) who showed that the part-list cue inhibition was absent when the part-list cues were aligned with the participants' subjective subcategories created during study.

Finally, the third and most crucial factor that produces an advantage for the control model is the effect of associative clustering. Raaijmakers and Shiffrin (1981) were able to show that the clusters sampled from the control models were richer in critical items that were not included in the part-list cues while the cued model's clusters typically included at least one cue word. They were able to show this by assuming that (1) the study items were clustered into associated groups of three, (2) that sampling a memory item from one of the clusters leads to recall of the sampled item and the other two items in the cluster immediately after, and (3) the control and cued models sample the same number of different triads during any set period of time. If these assumptions are held, it follows that the sampled clusters from the cued model contain fewer critical recall items because each cluster has at least one cue word while the sampled clusters from the control model will not always contain cue words and thus will contain more critical recall items. In this way, the retrieval structure of the models forces the cued model to recall more cue items than the control model. Because the model performance is compared on the total number of critical items recalled, the cued model ultimately has worse performance. This explanation is comparable to what the retrieval disruption hypothesis suggests is happening during part-list cuing. If the cued models are forced to retrieve fewer critical items because of the retrieval structure, then the part-list cues are disrupting the typical retrieval process of the model.

1.2 Adapting SAM to collaborative free recall of uncategorized lists

Given its success at modeling the part-set cuing effect in individual memory, adapting the SAM model to collaborative recall could provide valuable insights into the cognitive mechanisms behind collaborative inhibition. The SAM model could have been adapted to use either of the two prominent recall methods of collaborative recall tasks. The first possibility was the turn-taking method which was used by B. H. Basden et al. (1997) and Luhmann and Rajaram (2015). However, as mentioned previously, this method is less common in the literature, partly because the turn-taking method has been shown to increase memory intrusions (M. L. Meade & Roediger, 2009; Rajaram & Pereira-Pasarin, 2010). The second possibility was a free-for-all recall method in which participants recall freely whenever they retrieve a new word. This paradigm is the most popular within collaborative recall behavioral experiments and thus this recall method for cSAM was implemented.

The cSAM framework is defined as follows: a nominal group of N members has N models each carrying out free recall independently. A collaborative group of N members has N models storing associations independently (exactly as in the nominal group) but interacting during retrieval. The key assumption is that all N collaborative models use the most recent recall by any of the models as the next cue for retrieval, and all models increment associative strengths for any word recalled by the group. Note that these assumptions match the analogous assumptions for SAM applied to individual recall. All models in both groups use the same method for storing traces in memory, based on the SAM rehearsal buffer. Short-term memory is cleared before retrieval begins.

During retrieval for the nominal group, each of the N models retrieves independently in the normal SAM method that is summarized above. Figure 2.1 provides a flowchart of the free-for-all retrieval method for the collaborative group. There is a shared auditory buffer, called the group response, between the N models; which represents words spoken aloud by the models. The N models of the collaborative group are assumed to use each other's recalls as they occur, so it is necessary that their timing during retrieval be synchronized. Thus, all retrieval takes place in a sequence of fixed time steps. Each retrieval attempt takes one time step if no word is recovered after sampling. However, if a word is recovered it must be checked to see if the group has recalled it previously. This checking takes one time step. If the word has been recalled already, then a failure is counted. Whether a failure occurs in one time step or two, it is counted toward *Lmax* and *Kmax*. If the recovered word had not been previously recalled by a group member, then it is successfully recalled and every model in the group uses it as a probe cue for the next step of the search. This method can cause different models to accumulate different numbers of time steps as failures continue to mount, depending on whether a failure takes one or two time steps. Thus, the models can reach *Lmax* (and *Kmax*) at different times and at any given moment some models can be probing memory with context only while others are probing memory with word plus context.

Nonetheless, when any recall occurs, all models switch and use it to search memory. In addition, because the models can reach *Kmax* at slightly different times, models that have reached *Kmax* before other group members can continue to use retrieved items produced by other models as probes until *Lmax* is reached but cannot produce new items from context only search.

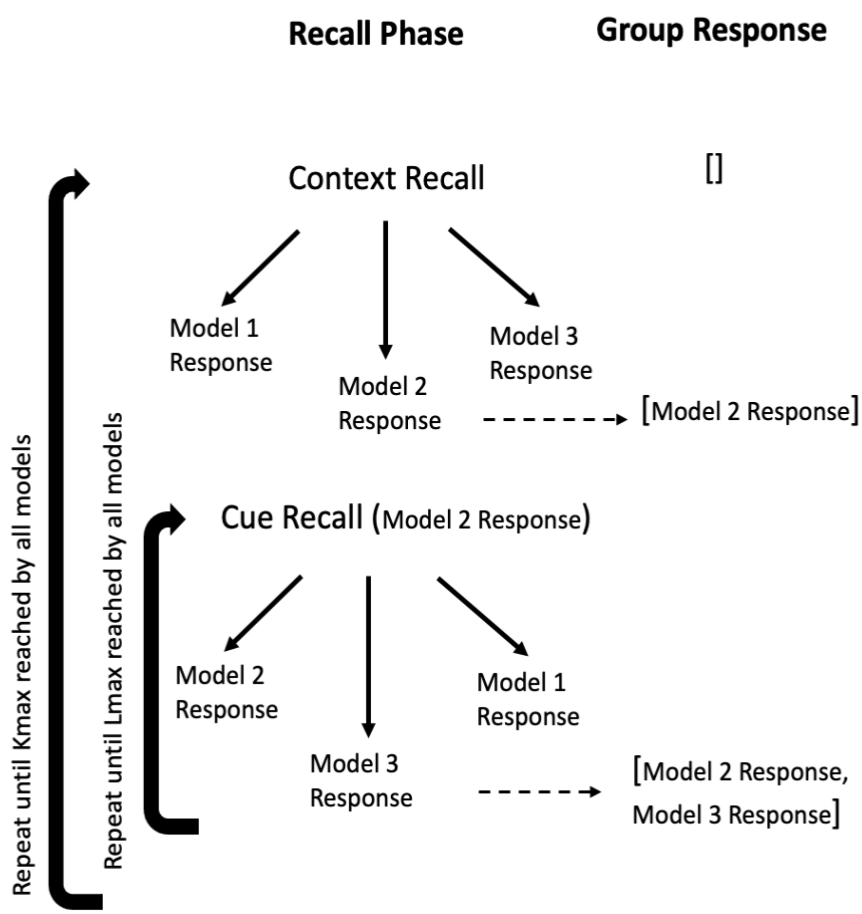

Figure 2.2. *Flowchart of collaborative recall between two or more SAM models. To begin recall, all models in the group start performing context recall. All models do this separately and the model with the fastest response (in this example Model 2's response was fastest) is added to the group response. Then, each model in the group uses the previous response (Model 2's response) to perform cued recall with the previous response as the cue word. Once again, the fastest model response is added to the group response (in this example Model 3's response was fastest). Cued recall continues until the stopping parameter Lmax is reached by all models in the group at which point the models begin context recall again. Once the stopping parameter Kmax is reached by all models recall ends.*

Thus, to summarize, the models for the nominal group (SAM) and the collaborative group (cSAM) are the same—only the experimental condition differs. The nominal group models use as a next cue the most recent word recalled; the collaborative group models use as a next cue the most recent word recalled by the group. The nominal group models increment associative strengths from cues used to traces of the current word recalled; the collaborative group models increment associative strengths from their cues used to traces of the current word recalled by the group.[2]

1.3 Adapting SAM to collaborative free recall of categorized lists

Basden et al., (1997) explored collaborative groups free recalling categorized lists. Collaborative inhibition was found, and the magnitude was larger for 6 categories of size 15 than 15 categories of size 6. To simulate their design, all items in all categories are presented to each individual for study in mixed order, and retrieval begins after short-term memory is cleared. Nominal and collaborative groups are assumed to have three members. The models applied are very similar to those for uncategorized lists with a few slight modifications. During storage, the buffer model is applied as usual, except that the total association between two words in the same category is larger than in the uncategorized model by 0.25, and the association between two words in different categories is larger than in the uncategorized model by 0.005. Retrieval is the same as for the uncategorized model (except that the various associative strengths will be different, due to the category structure of the lists).

[2] Some group members continue recall beyond the point they would have done (at *Kmax*) because they are induced to continue search by a recall by some other group member that has not yet reached *Kmax*. This factor helps increase group recall compared to nominal (to a tiny degree). Collaborative inhibition occurs regardless.

2. Modeling Results

2.1 Parameter estimation of individual data

Before attempting to fit cSAM to aggregate recall data, 3 parameters for the nominal and collaborative models (*e, f,* and *g*) were fit using individual data (from Choi, Blumen, Congleton, & Rajaram, 2014) to investigate any informative parameter differences between nominal and collaborative groups. The data provided consisted of recalls from a list of 50 uncategorized words for both nominal and collaborative groups. There were 36 groups represented in this data set, 18 nominal and 18 collaborative. To estimate these parameters using individual level data, the forest.minimize optimization function from the scikit-optimize Python library was used. This optimization function begins by modeling a function, in this case the retrieval process of cSAM, using a decision tree-based regression model. The model is then improved by evaluating the function at the next best point—thus finding the minimum value of the function in the smallest number of steps.

A Kolmogorov-Smirnov test for goodness-of-fit was performed on each parameter distribution to check for normality. All were found to be significantly non-normal, leading us to employ a nonparametric test to compare collaborative and nominal distributions for each parameter. Table 2.2 displays the means, standard deviations, and medians for best-fitting parameter values across the participants. There were no significant differences in the median parameters for the collaborative and nominal groups, suggesting that CI emerges from the structure of the model when placed in this experimental paradigm and not the values of the parameters. Thus, the performance of the collaborative and nominal groups while keeping all parameters constant between groups was compared.

Table 2.2. *Descriptive Statistics of Parameter Distributions*

Parameter		Collaborative	Nominal	Mood's Test p-value
e	Mean (sd):	0.72 (.14)	0.71 (.14)	0.11
	Median:	0.70	0.7	
f	Mean (sd):	0.73 (.14)	0.74 (.14)	0.42
	Median:	0.71	0.71	
g	Mean (sd):	0.73 (.14)	0.73 (.15)	0.74
	Median:	0.70	0.7	

2.2 Fitting cSAM to aggregate collaborative recall data

SAM and cSAM were then used to predict nominal and collaborative groups of three members freely recalling 40 word lists of uncategorized words (Weldon & Bellinger, 1997) and the major findings by B. H. Basden et al. (1997) for groups of three members freely recalling categorized lists of two types (B. H. Basden et al., 1997). Because no significant parameter differences emerged in the previous section, SAM and cSAM were set to use the same fixed parameter values (Table 2.1). The only difference in the models lies in the fact that the collaborative group members use the most recent recall by any group member for their next searches of memory while the nominal group members use their own most recent recall for their next searches of memory.

When all the values were set to those used in Raaijmakers and Shiffrin (1981) for part list cuing, CI emerged. However, the levels of recall predicted for uncategorized and categorized word lists were slightly off target. Thus the values of a, b, and c were slightly changed (by hand) to .08 to match the observed data for uncategorized lists from (Weldon & Bellinger, 1997), and slightly changed to .07 to match the observed data for categorized lists from B. H. Basden et al., (1997). The values of h and i were not needed to predict CI but were set (by hand) so that the predictions for categorized lists would better match the observed data. The other parameter values for uncategorized lists were those used in the 1981 article. The results of our parameter estimation show that parameter estimation is ultimately unnecessary as no significant differences emerge and CI is predicted for almost any parameter values.

Figure 2.2 shows the results of fitting the model to individual, nominal, and collaborative recall of uncategorized lists, using the first 11 parameters shown in Table 1. Figures 2.3 and 2.4 show the results of fitting the model to individual, nominal, and collaborative recall of categorized lists, using all 13 parameters shown in Table 2.1.

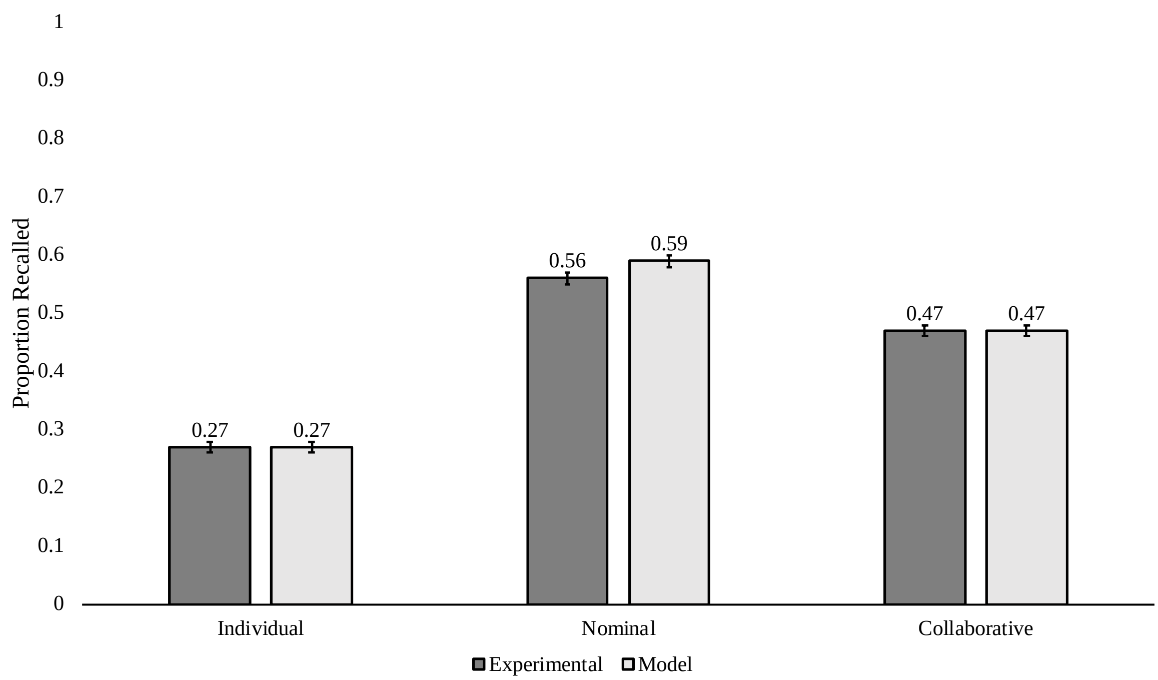

Figure 3.2. *SAM model fit to uncategorized list data taken from the original Weldon and Bellinger (1997) paper detailing collaborative inhibition. Subjects were tested in groups of 3 on a list of 40 unrelated words.*

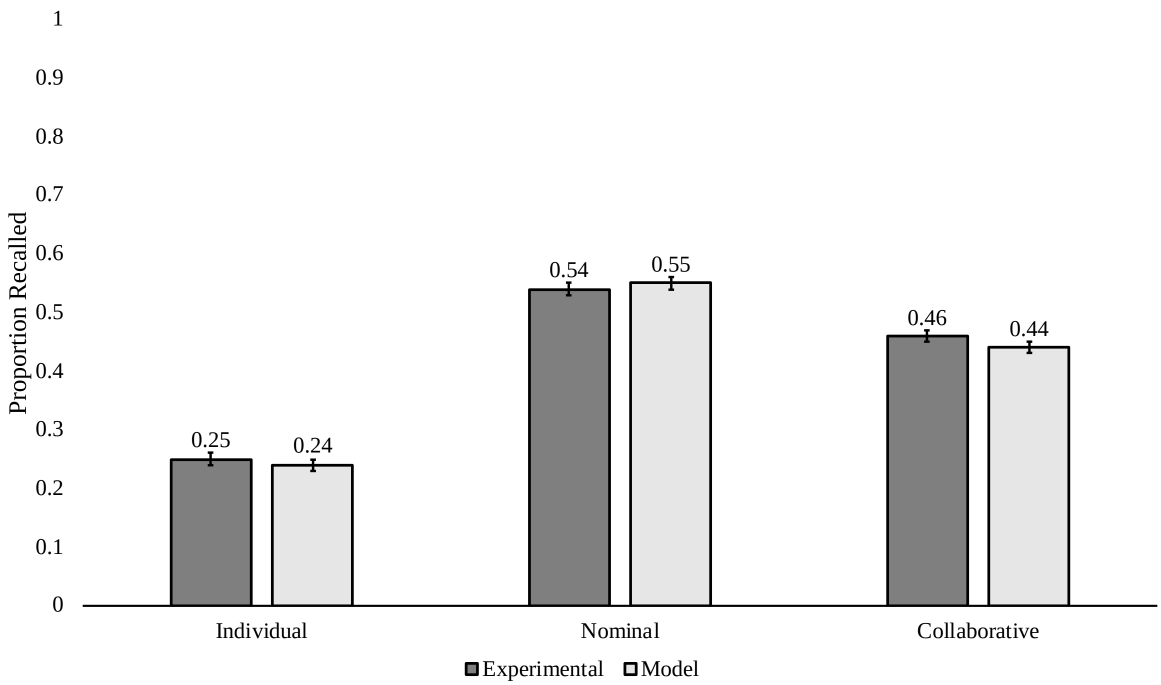

Figure 2.4. *cSAM model fit to categorized list data from B. H. Basden et al. (1997). Subjects in groups of 3 were asked to recall from a list of 90 words grouped into 6 total categories with 15 items in each category. The larger category size results in a more prominent collaborative inhibition effect.*

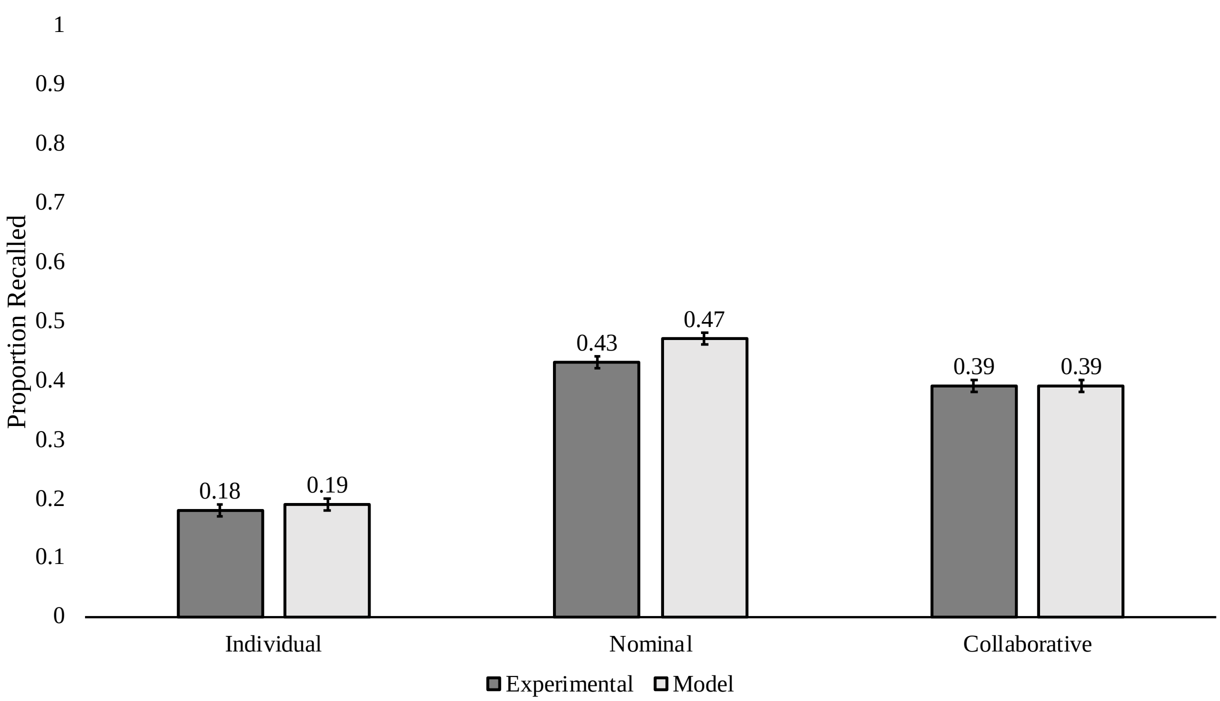

Figure 2.5. *cSAM model fit to categorized list data from B. H. Basden et al. (1997). Subjects in groups of 3 were asked to recall from a list of 90 words grouped into 15 total categories with 6 items in each category. The smaller category size results in a less prominent collaborative inhibition effect.*

The experimental data from B. H. Basden et al. (1997) (used to fit the model in Figures 2.3 and 2.4) supports the retrieval disruption hypothesis because collaborative inhibition is stronger when study materials are less organized. In the first condition (Figure 2.3) study materials are less organized because the group sizes are larger, allowing more room for idiosyncratic organization within categories. In the second condition (Figure 2.4) the study materials are more organized because the group sizes are smaller, allowing less room for idiosyncratic organization within categories. When the internal organization of study items is dissimilar between group members, collaborative inhibition increases because the cues from other group members are more likely to disrupt individual search strategies. Thus, Figures 2.2, 2.3, and 2.4 show that the SAM/cSAM modeling is capable of naturally reproducing the key findings of collaborative inhibition.

3. Discussion

The implications of collaborative memory research are much larger than participants recalling lists of words in experimental settings. The cognitive mechanisms being studied by this basic science are the same that play a role in crucial applied phenomena such as the spread of misinformation, memory contagion, fake news, eyewitness testimony, and even conspiracy theories. To date, there are no formal computational frameworks within which to understand how the memory mechanisms of individuals interact to produce emergent phenomena when collaborating. In this chapter, a first step towards this goal was achieved by modifying the well-validated SAM model of Raaijmakers and Shiffrin (1981), and providing an existence proof that the collaborative SAM framework can produce the basic patterns of collaborative inhibition seen in experimental data. In addition to basic uncategorized lists, cSAM naturally produces the patterns seen in categorized lists, namely greater collaborative inhibition when study materials are less organized. Importantly, each SAM model in isolation would still retain the explanatory power for the range of behavioral phenomena in individual memory paradigms, providing a unified model to understand both individual and collaborative memory.

When fitting the three learning parameters, there were no significant differences between nominal and collaborative groups. No significant differences in optimal parameter values encouraged a comparison of the performance of the collaborative and nominal groups while keeping all model parameters the same between both groups and found that collaborative inhibition persisted. These results indicate that collaborative inhibition is being caused by a mechanistic or structural difference in cSAM that is not captured by parameter differences.

One such mechanism is likely the same mechanism responsible for part-list cuing deficits shown in SAM. SAM produces part-list cuing deficits when a retriever uses each provided word as the cue for the next memory search. Similarly, cSAM predicts collaborative inhibition when the retrievers use the most recent recalled word by any group member as a cue for the next memory search. In both cases, the cue words being used are ineffective compared to the ones used individually because the cues are mismatched to the subjective organization formed by the individual during list study. In the case of part-list cuing, the experimenter might randomly provide cues that are strongly associated to each other and poorly associated to the words that are the object of retrieval. In the case of collaboration inhibition, when group member A recalls W_1 it could be a good cue for A because A has stored a strongly interconnected group [W_1,W_2,W_3] and none of these words has yet been recalled. But group members B and C might have nothing in their memory that is strongly connected to W_1 other than words previously recalled by the group. Thus, B and C would have their retrieval disrupted when they use W_1 that was produced by A.

The simulations show that the net effect of these retrieval disruption factors, instantiated in cSAM as induced use of ineffective cues, produces collaborative inhibition. In addition, Figures 2.3 and 2.4 show that the amount of predicted collaborative inhibition is slightly larger for fewer but larger categories, as observed by Basden et al. (1997) and used to argue for retrieval disruption. That is, when the categories are smaller, there is less opportunity for idiosyncratic organization within each category, and hence less opportunity for different subjective organizations by different group members. It should be noted that the differences in subjective organization in cSAM are produced by differences in the way different models rehearse during study, because rehearsal in SAM and cSAM is a stochastic process. There are of course numerous other reasons why subjective organization might differ among group members, so it is interesting that the limited degree of subjective organization produced by stochastic storage is sufficient to produce significant collaborative inhibition.

The simulations in this chapter suggest that collaborative inhibition arises from a mismatch of the cues used to search memory and the differing subjective organizations of the group members. The SAM model has been shown to capture the primary processes of recall by individuals, and the present results go further and show the same processes of recall and memory search provide a unified model to understand both individual and collaborative memory. Given that cSAM has been shown to capture the standard patterns of collaborative inhibition seen in the literature, it would be natural to extend it to make predictions and suggestions for future experimental studies. Such studies might help us understand not just the factors operating to produce collaborative inhibition, but the factors critical for searching memory. The use of a formal computational framework helps differentiate between theories of group memory that are currently unresolvable using only experimental data and in addition helps the field generate new predictions and experiments to advance the study of memory storage and retrieval.

CHAPTER 3

APPLICTIONS OF cSAM: MECHANISTIC EXPLANATIONS AND PREDICTIONS FOR COLLABORATIVE INHIBITION

In the previous chapter, the collaborative Search of Associative Memory (cSAM) framework was introduced and validated on common experimental findings within the collaborative memory field. The presence of collaborative inhibition in cSAM simulations when collaborative and nominal groups have the same parameter settings suggests that there is a mechanistic or structural difference between models that causes collaborative inhibition. Additionally, the patterns produced during the categorized list simulations suggests that retrieval disruption may be a contributing factor to collaborative inhibition in cSAM. As of now, retrieval disruption is the only mechanistic difference that has been considered between collaborative and nominal model groups, however, there are several other factors that may affect collaborative inhibition in cSAM simulations. Memory homogenization, shared background knowledge (expertise), group size, and post-collaborative effects have all been shown to moderate collaborative inhibition and provide insight into additional mechanistic group differences (Barber et al., 2015; Luhmann & Rajaram, 2015; M. L. Meade et al., 2009; Thorley & Dewhurst, 2007).

Investigating these areas will play a critical role in gaining a deeper understanding of what is driving collaborative inhibition in cSAM while also making predictions about understudied effects in the experimental literature.

1. Memory Homogenization

Before the cSAM framework, the only other attempt at modeling collaborative memory was a verification step of a study looking at information transmission in networks using an agent-based modeling approach (Luhmann & Rajaram, 2015). Though the main goal of their study was not to model collaborative inhibition in episodic recall tasks, collaborative inhibition was nevertheless observed in the recall of groups of 3 agents. Additionally, the model introduced in this study was able to generate some predictions of understudied effects within the collaborative memory field, namely the effect of group size on collaborative inhibition. While this model included psychologically based agents that were able to encode and retrieve memories, the implementation was not as mechanistically extensive as cSAM and may have produced collaborative inhibition in different ways.

While verifying their agent-based model, Luhmann and Rajaram (2015) found evidence of collaborative inhibition. However, their explanation for why collaborative inhibition occurred in their model was not due to retrieval disruption but rather by the agents' memories homogenizing as they collaborated. They explain that after the study phase of the collaborative recall task, the agents each had an idiosyncratic activation pattern over the study items. Learning during the collaborative recall task decreases the diversity of the agents' memory representations, which the authors believe reduced collaborative recall performance and caused collaborative inhibition. While the agent-based model was able to successfully induce collaborative inhibition, it was likely that the reason was not due to retrieval disruption. Additionally, this explanation seems to contradict predictions of the retrieval disruption hypothesis. According to the hypothesis, if group members memories are more similar, and their retrieval strategies are more similar, then external cues provided by group members should not disrupt retrieval nearly as much. Several studies have shown that collaborative encoding, which causes more similar retrieval organization, reduces collaborative inhibition (B. H. Basden et al., 1997; Finlay et al., 2000). Thus, it is bold to claim that memory homogenization, which predicts opposite outcomes of memory similarity, is the sole cause of collaborative inhibition. However, it is possible that there is a dual effect of both memory homogenization and retrieval disruption at play in both experimental and modeling studies. Unfortunately, it is impossible to isolate memory homogenization in a behavioral study, thus this possibility has never been proposed before.

The goal of this section is to determine whether memory homogenization could be a cause of collaborative inhibition in cSAM. Previously, cSAM has shown support for the retrieval disruption hypothesis, but the results of the Luhmann and Rajaram (2015) study suggest there may be other causes to consider that can't be studied in behavioral research. To tease apart the underlying cause of collaborative inhibition in cSAM, memory homogenization during collaborative recall was stopped by preventing learning during the retrieval phase. If collaborative inhibition is still present when model memories remain diverse, then memory homogenization can be ruled out as the sole cause of collaborative inhibition in cSAM.

1.1 Learning allowed during collaborative retrieval

Both the original individual SAM models and the cSAM models learn during retrieval. Learning in cSAM models is the same as individual models and is controlled by the *e*, *f*, and *g* incrementing parameters (see Figure 3.1). While learning does occur during collaborative retrieval via the incrementing parameters, it has not yet been shown that model memories homogenize when retrieving in groups. To measure the similarity of model memories during retrieval, the overall cosine similarity between each model's word association memories was recorded. The cosine similarity between two vectors is their dot product divided by the product of their magnitudes.

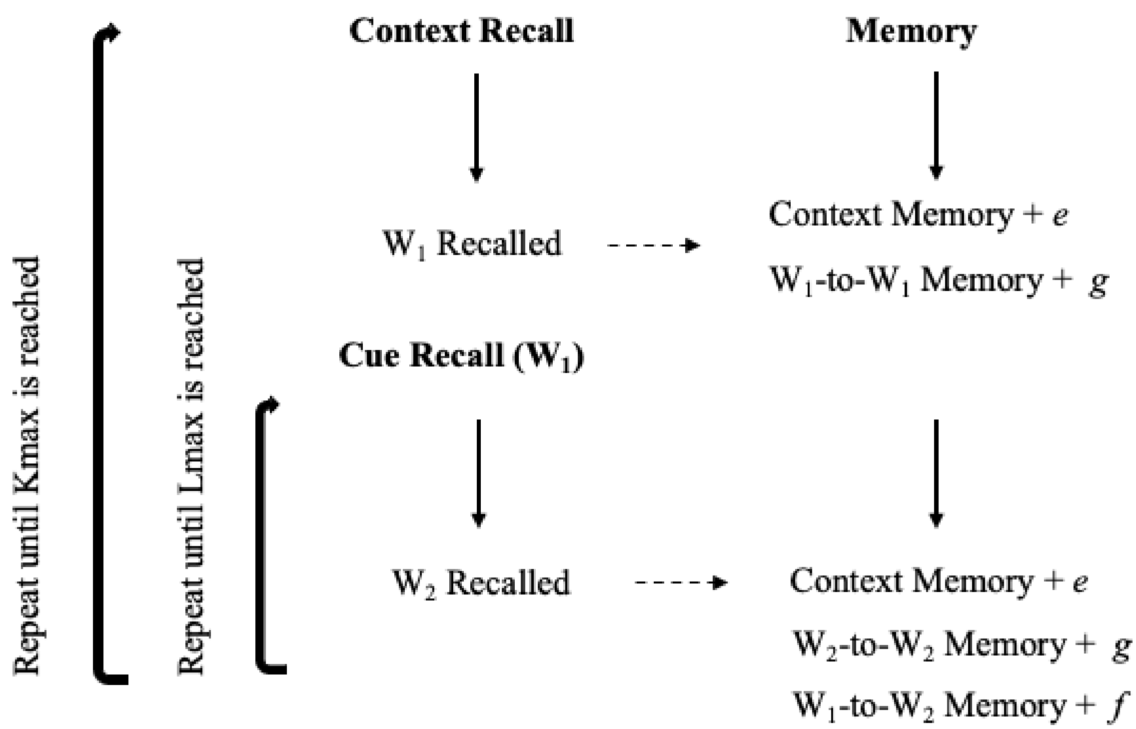

Figure 3.6. *Diagram of Learning during Retrieval in SAM*

To measure the change in cosine similarity between multiple model memories, the individual word representations within each model at 20 different timestamps during retrieval were compared. The timestamps used in these calculations are the points during retrieval when a word is successfully recalled by the group and the model memories are updated (see Figure 3.1). This process works as follows: if there are 2 study words on the study list and 2 models in a group, then the mean cosine similarity of word 1 (represented by Row 1 in Figure 3.2) between models 1 and 2 and word 2 (represented by Row 2 in Figure 3.2) between models 1 and 2 for timestamp 1 is calculated and averaged. This process is repeated for all 20 timestamps. For a visualization of this process over one timestamp, see Figure 3.2.

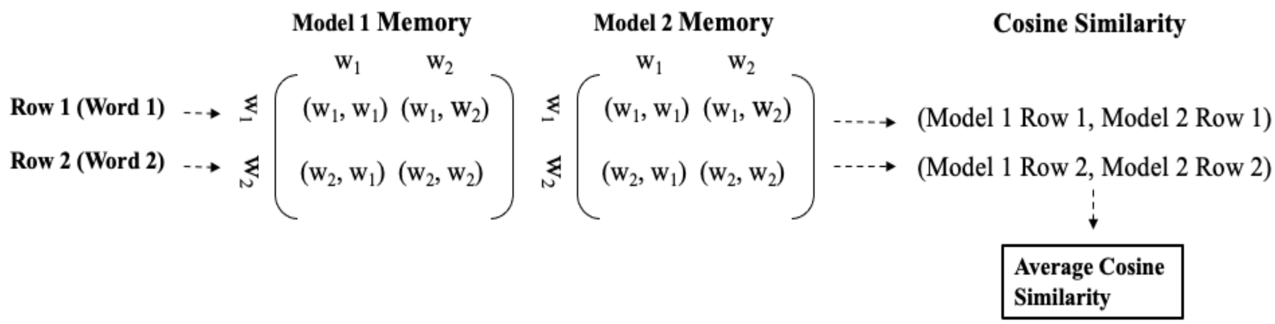

Figure 3.7. *Cosine similarity calculation during one timestamp in collaborative retrieval. Row 1 represents the association vectors for word 1 in Model 1 and Model 2, respectively. The cosine similarity of these two vectors is calculated and recorded. Row 2 represents the association vectors for word 2 in Model 1 and Model 2, respectively. Again, the cosine similarity between these two vectors is calculated and recorded. Then, the mean of these two cosine similarities is calculated and recorded. This process is repeated for each timestamp during collaborative retrieval.*

The result of this process is Figure 3.3 which shows the average cosine similarity (over 100 collaborative retrieval simulations) of the word association memories of a model group over 20 timestamps during retrieval. During retrieval, the association memories of models in a collaborative group do become more similar to each other (homogenization), whereas models in the nominal group experience no memory homogenization. This is consistent with the finding that model memories homogenize during collaboration and become less diverse over the course of collaborative retrieval (Luhmann & Rajaram, 2015). Additionally, model memories became significantly more similar to each other during collaborative recall ($M = 0.58$) than during nominal recall ($M = 0.49$), $t(198) = 27.24$, $p < 0.001$.

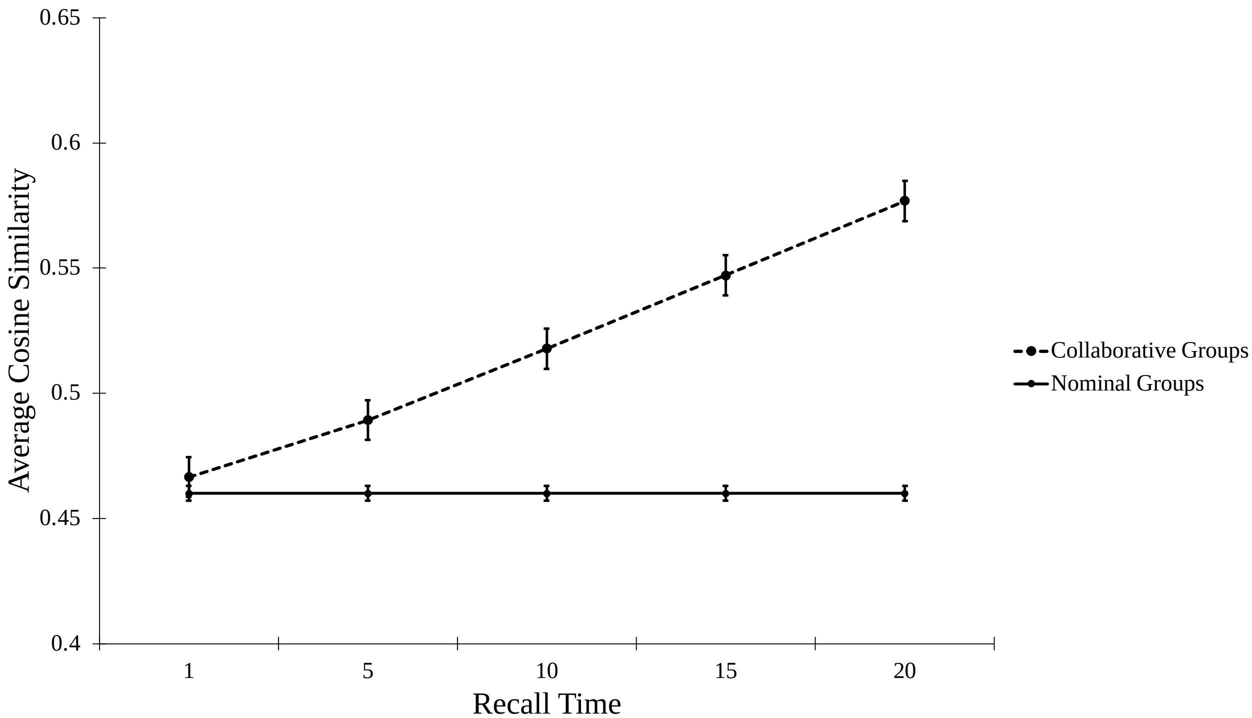

Figure 3.8. *The mean cosine similarity (over 100 collaborative retrieval simulations) of cSAM models' associative memories during collaborative retrieval over 20 timestamps when learning is allowed during retrieval. Collaborative group memories become more similar to each other while nominal group memories do not.*

The results shown in Figure 3.3 demonstrate that memory homogenization occurs in cSAM models over the course of retrieval. To determine whether memory homogenization is the sole cause of collaborative inhibition in cSAM, memory homogenization will be prevented. If collaborative inhibition is still present, this will suggest that an alternative mechanism, such as retrieval disruption, contributes to collaborative inhibition.

1.2 Learning prevented during collaborative retrieval

To determine whether memory homogenization is the sole cause of collaborative inhibition, learning during retrieval was prevented—something that is impossible during a behavioral experiment. If no learning occurs during retrieval, then the model memories will not homogenize and become less diverse. If collaborative inhibition persists, this will support the claim that memory homogenization is not the sole cause of collaborative inhibition in cSAM.

To accomplish this task, the parameters responsible for learning during recall in the collaborative models (*e*, *f*, and *g*) were set to 0 so that the models do not learn at all during retrieval. This should prevent the model memories from homogenizing over retrieval. To ensure the manipulation had the desired effect on the memory structures, the method of measuring model memory similarity over retrieval was repeated. Figure 3.4 shows the change in model memory similarity when learning was prevented. When learning is prevented, model memories do not homogenize over the course of retrieval. In fact, model memories stay exactly as similar to each other as they were at the beginning of retrieval.

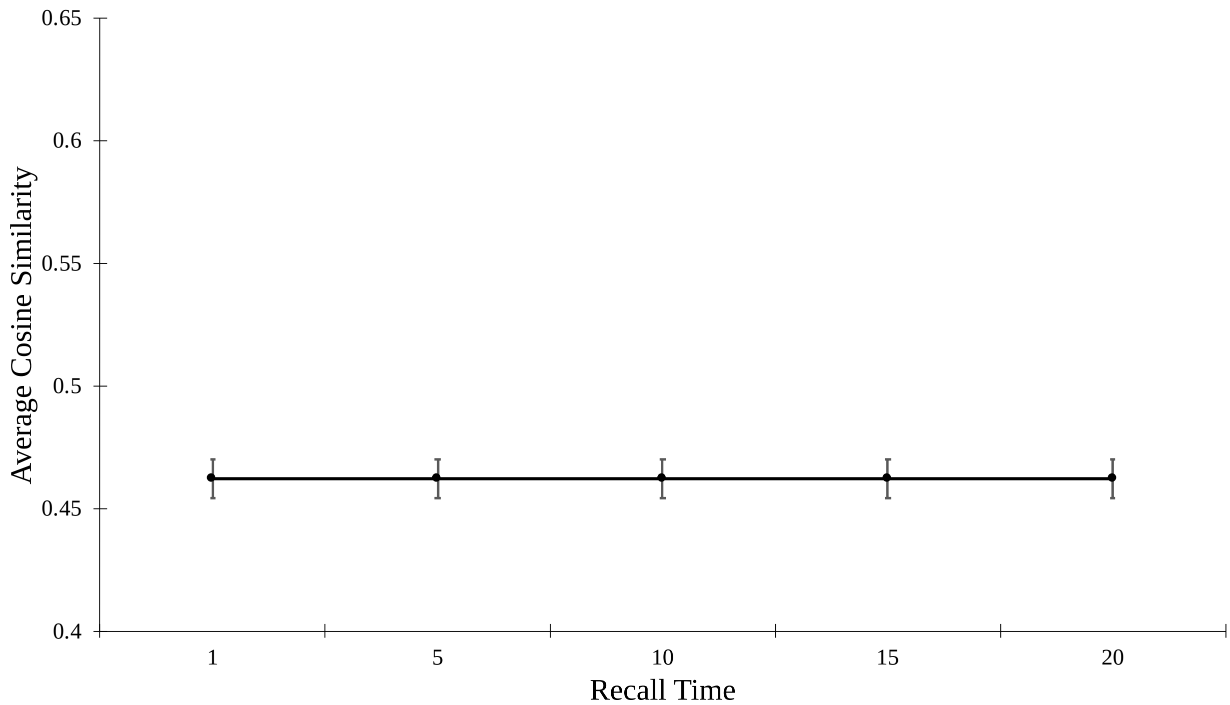

Figure 3.9. *The mean cosine similarity (over 100 collaborative retrieval simulations) of cSAM models' associative memories during collaborative retrieval over 20 timestamps when learning was prevented. Collaborative group memories no longer become more similar during retrieval.*

After validating that the prevention of learning during retrieval also stopped memory homogenization between models, the effect of preventing learning on collaborative inhibition was evaluated. In Figure 3.5, the amount of collaborative inhibition in the learning allowed and learning prevented conditions of collaborative retrieval was compared. Figure 3.5 shows that collaborative inhibition persists in the learning prevented condition, $t(198) = 8.34$, $p < 0.001$, where the model memories do not homogenize over retrieval, however, the size of the inhibitory effect is diminished. This suggests that collaborative inhibition in cSAM simulations is not caused solely by increased memory homogeneity over retrieval.

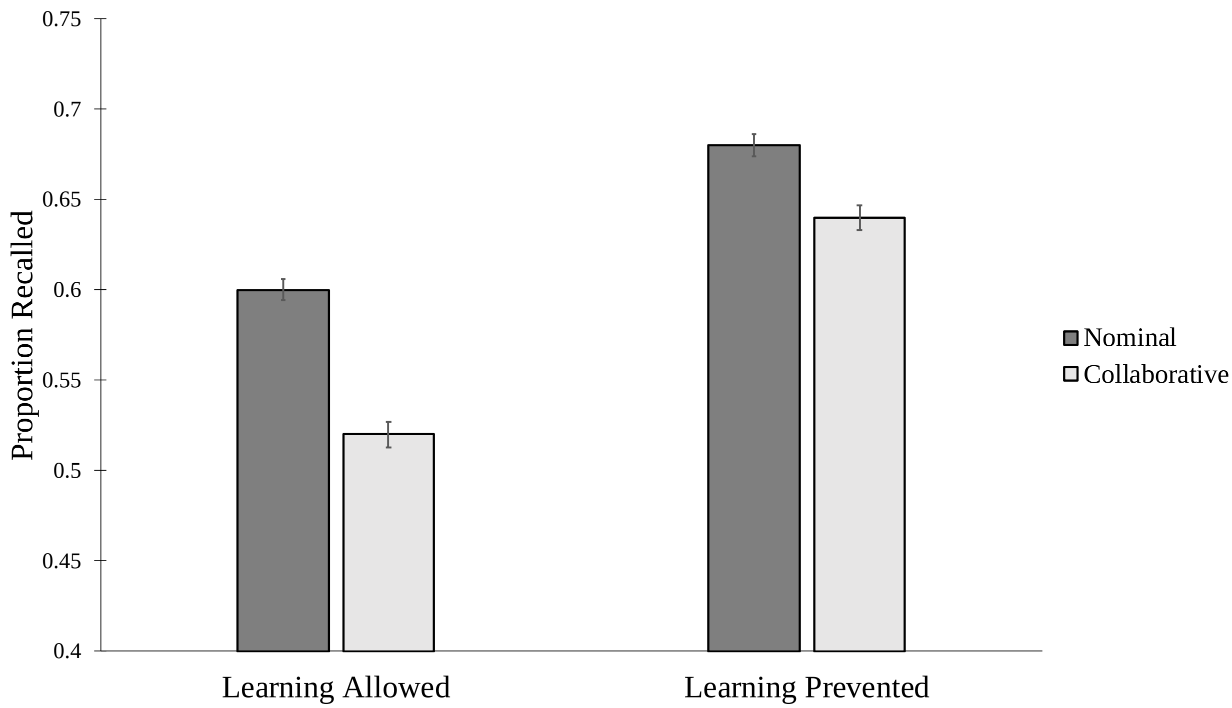

Figure 3.10. *Comparison of collaborative inhibition in learning allowed and learning prevented conditions of collaborative retrieval.*

In Figure 3.5, collaborative inhibition was still present in the case where model memories did not homogenize. However, the overall proportion recalled increased in the learning prevented condition and the effect size of collaboration was diminished. The increase in proportion recalled is a natural effect of how learning works in cSAM. During the word cue recall phase of retrieval, learning causes words that are recalled to have a higher association with the cue word. A higher association between words means that those words are likely to be recalled together. By the end of retrieval, due to the changing memory association matrix, it is more difficult for models to produce words that haven't already been recalled. Eventually, the models can't produce any new words and retrieval ends (when *Kmax* is reached). However, when learning is prevented during retrieval, the models' memories are not updated when a word is recalled using a word cue. Consequently, it is less difficult for models to produce new words and the models can continue retrieval for longer, producing more words overall. This also plays into the effect memory homogenization has on collaborative inhibition. When learning is not prevented, model memories in collaborative groups become more similar over recall than model memories in nominal groups. This leads to the collaborative groups having more difficulty recalling new words, thus increasing collaborative inhibition.

To determine whether this explanation has merit, the average number of words recalled from a list of 50 unrelated words in the learning allowed and learning prevented conditions were compared. Groups of models in the learning prevented condition recalled significantly more words on average than groups in the learning allowed condition $t(198) = 12.17$, $p < 0.001$ (see Table 3.1). This result supports the previous explanation for a higher overall proportion recalled in the learning prevented condition.

Table 3.3. *Average number of words recalled over 100 simulations of collaborative retrieval*

	Mean (std)
Learning Allowed	26.6 (3.2)
Learning Prevented	32.7 (3.3)

1.3 Discussion

The goal of part 1 of this study was to determine how model memories change during collaborative retrieval. Model memories were found to naturally homogenize in a collaborative setting (see Figure 3.3) and models in collaborative groups experienced more memory homogenization than those in nominal groups. The goal of part 2 was to determine whether decreased memory diversity during retrieval is solely responsible for collaborative inhibition. Preventing learning during retrieval also prevented memory homogenization and when learning was prevented, collaborative inhibition persisted, though with a diminished size of effect (Figure

3.5). The results of this section suggest that while memory homogenization is not the only cause of collaborative inhibition in cSAM, it does still contribute to the inhibitory effect. This is an entirely novel finding as memory homogenization is impossible to isolate as a contributing factor in behavioral experiments. This finding further supports the idea that retrieval disruption is not the only mechanism responsible for collaborative inhibition in cSAM.

While the results of this study do not support the claim that collaborative inhibition is caused only by memory homogenization, the idea of shared memories after collaboration is not unsupported by the literature. Blumen and Rajaram (2008) found that after collaborative recall, participants have an increase in overlap of their post-collaborative individual recall—suggesting that group members memories do homogenize. Additionally, Congleton and Rajaram (2014) found that the presence of collaborative inhibition may be responsible for shared group memories that arise after collaborative recall. They found that as the size of the collaborative inhibition effect increases, so does the amount of shared memory organization and shared memories after collaboration. It is believed that when group members' retrieval strategies are disrupted, they are more likely to adopt the organization created by the group for subsequent instances of recall instead of continuing to use their original individual organizations.

The findings from these experimental studies suggest that there is a relationship between collaborative recall and subsequent shared memories that may arise due to memory homogenization during collaborative recall. CSAM shows the same pattern of memory homogenization over retrieval (Figure 3.3) and supports these experimental findings. However, while memory homogenization is present during collaborative retrieval, it does not eliminate collaborative inhibition when prevented—suggesting other mechanisms, such as retrieval disruption, may be involved.

2. Shared Background Knowledge

Researchers interested in the effect of relationships on collaborative inhibition have shown that pairs of friends and married couples both show a reduction in collaborative inhibition compared to stranger pairs (Andersson & Ronnberg, 1995, 1996; Johansson et al., 2000). To date, there has been one study investigating the effect of expertise on collaborative inhibition. M. L. Meade et al. (2009) found that if subjects had expertise in the area of study, collaboration was beneficial compared to nominal groups. They compared the performance of expert pilots, novice pilots, and non-pilots on a collaborative memory task involving aviation scenarios and found that while the non-experts still suffered from collaborative inhibition, the expert pilots benefited from collaboration during recall. All studies that investigate the effect of shared background knowledge have found that those with more shared background knowledge are less affected by collaborative inhibition. However, it is unclear whether this effect is due to shared background knowledge or communication proficiency with group members.

It is impossible to test tease apart the effect of shared background knowledge and communication when testing groups of friends or married couples as both shared background knowledge and some level of communication proficiency are natural consequences of building and maintaining relationships. Studies using experts may be a way to separate the effects of shared background knowledge and communication proficiency on collaborative inhibition, however, aviation experts are not appropriate for this goal. M. L. Meade et al. (2009) mention that aviation expertise includes in depth training in communication skills which may be responsible for the faciliatory effect they found. It is uncertain whether collaborative facilitation in experts is unique to certain domains of expertise that involve communication training or if this result is generalizable to all experts. Recently, C. B. Harris, Barnier, Sutton, and Savage (2019) investigated the relevant communication patterns present when collaboration between married couples is beneficial and found that some behaviors, such as cuing of group members and successful responses to those cues indicates a high level of group coordination and is predictive of a faciliatory effect of collaboration. It is possible that the pilots observed were using similar communication patterns, allowing them to act with a high level of coordination during collaborative recall. Unfortunately, there is no available research comparing different forms of expertise to address this question. The goal of this section is to determine whether shared background knowledge reduces or reverses collaborative inhibition or if communication may play a larger role in previous findings than originally thought.

2.1 Simulating shared background knowledge in cSAM

To tease apart the role of shared background knowledge and communication in reversing or mitigating collaborative inhibition, shared background knowledge in cSAM was simulated while communication between the models during recall remained the unchanged. To simulate shared background knowledge in cSAM, the model initiation method was altered. In the original cSAM models, as soon as models are initiated, they create two empty memory stores: the context association vector and the word association matrix. The models then encode the words on the study list and fill in the empty memory stores according to the *a*, *b*, *c*, and *d* parameters. This initiation process was changed by randomly filling in the word association matrix with values pulled from a log normal distribution before the models encode any study items (see Figure 3.6). The log normal distribution is a probability distribution of a random variable whose logarithm is normally distributed. This distribution was used because it does not allow negative values like the normal distribution as having a negative association between study items is not supported by cSAM. By altering the standard deviation of the log normal distribution, the similarity of model background knowledge was manipulated before encoding. To measure the similarity of the model memories, the average cosine similarity of the word association matrix between all group members was calculated (see the process outlined in Figure 3.2 in section 1 of this chapter).

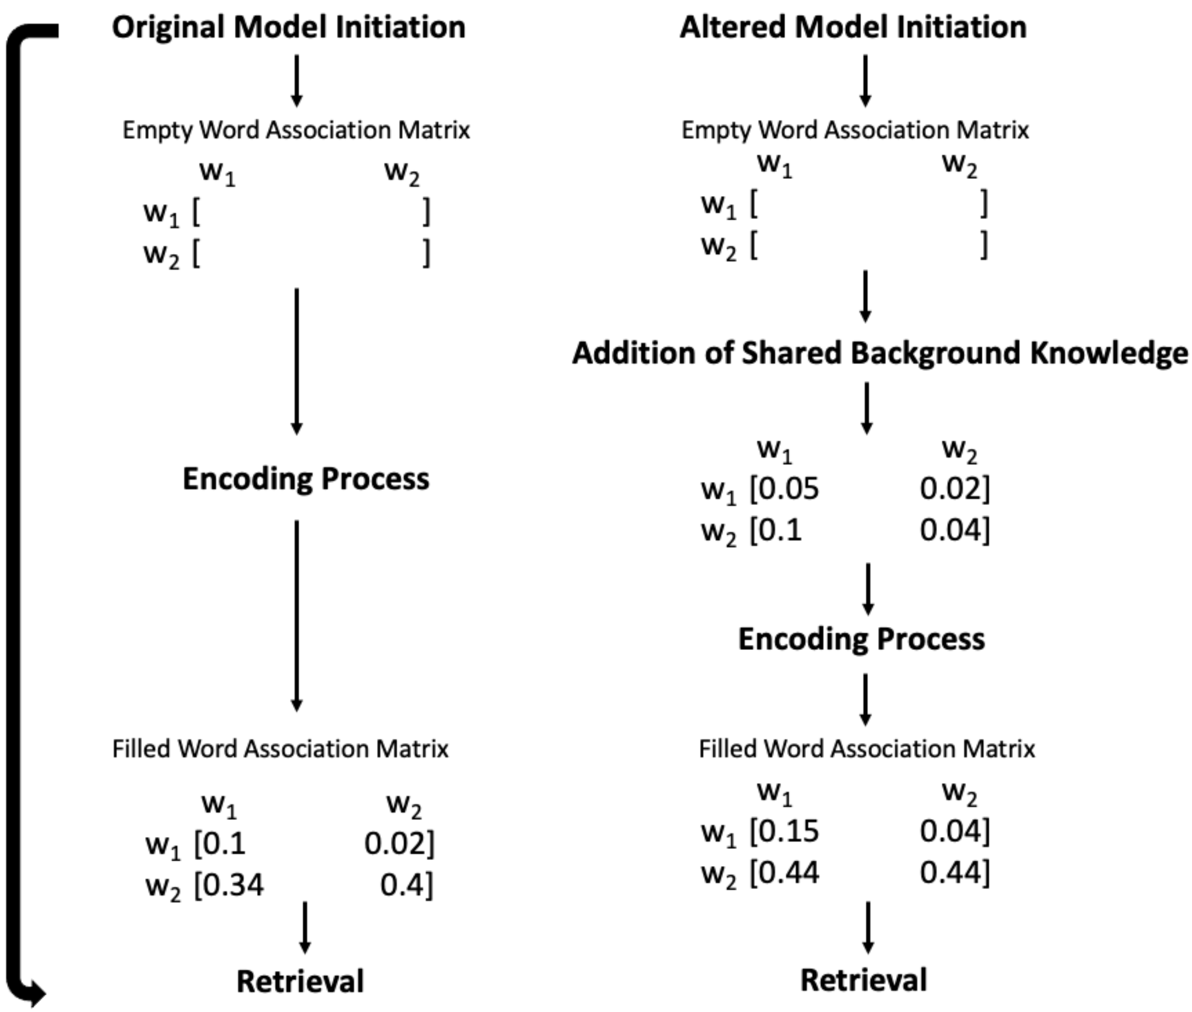

Figure 3.11. *Process of altering model initiation to include shared background knowledge before encoding process.*

2.2 Results

The effect of shared background knowledge on group performance was determined by comparing the proportion recalled by nominal and collaborative groups as background similarity between models increased before encoding. Nominal and collaborative model parameters were kept constant between groups for the simulations. Linear regression was used to determine if background similarity significantly predicted the proportion recalled by both groups. For the nominal groups, the fitted regression model was: Proportion Recalled = -0.18(Background Similarity) + 0.79. The overall regression was statistically significant ($R^2 = 0.82$, $F(1, 30) = 132.24$, $p < 0.001$) indicating that background similarity predicts the overall proportion recalled by the nominal groups. For the collaborative groups, the fitted regression model was: Proportion Recalled = -0.17(Background Similarity) + 0.69. The overall regression was statistically significant ($R^2 = 0.81$, $F(1, 30) = 117.77$, $p < 0.001$) indicating that background similarity predicts the overall proportion recalled by the collaborative groups. Figure 3.7 shows the relationship between shared background similarity and the performance of nominal and collaborative groups.

A between-participants two-way ANOVA was used to test the main and interaction effects of group type (collaborative vs. nominal) and similarity level (high vs. low) on the average proportion recalled. There were significant main effects of group type, $F(1, 196) = 7728.17$, $p < .001$, $\eta^2 = 0.98$ and similarity level, $F(1, 196) = 27076.79$, $p < .001$, $\eta^2 = 0.99$, both with large effect sizes. And, there was a statistically significant interaction between group type and similarity level, $F(1, 196) = 20.78$, $p < .001$, $\eta^2 = 0.10$, with a medium effect size. Figure 3.8 shows the comparison between low and high similarity starting background and its effect on proportion recalled for both group types.

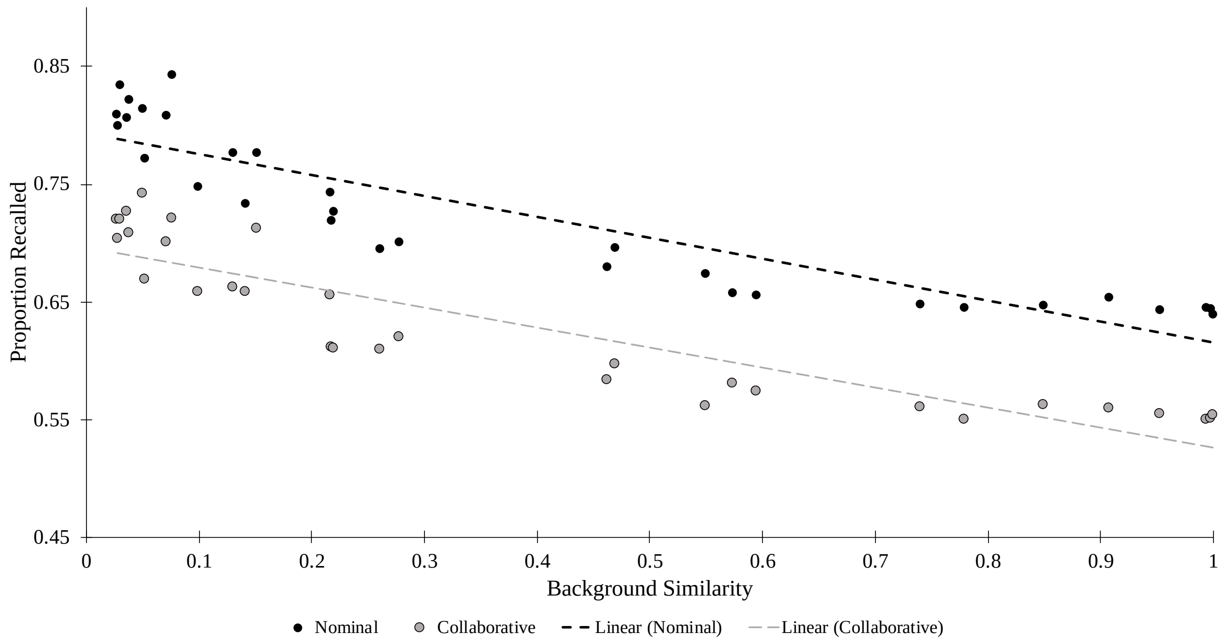

Figure 3.12. *The relationship between shared background similarity (cosine similarity) and the proportion recalled by nominal and collaborative groups.*

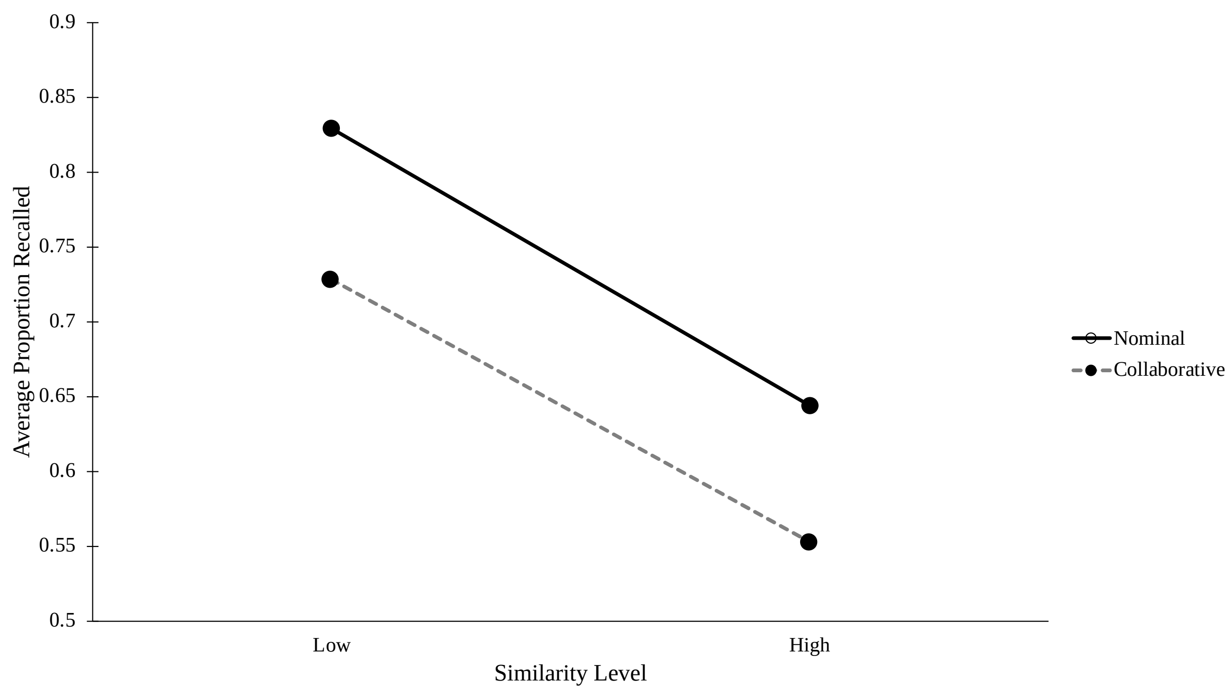

Figure 3.13. *The interaction between similarity level (high vs. low) and recall method (collaborative vs. nominal) on average proportion recalled on a list of length 50.*

Pairwise comparisons in Table 3.2 show that collaborative inhibition is present in both low and high background similarity conditions. Additionally, there is more collaborative inhibition in the low background similarity condition, $t(98) = 4.92$, $p < 0.001$. This effect is also shown in Figure 3.9 for visual clarity.

Table 3.4. *Mean differences between nominal and collaborative groups in high and low background similarity levels.*

Similarity Level	Group Type	Mean Difference	*t*	*p*	*d*
Low	Nominal- Collaborative	.10	71.76	<0.001	14.28
High	Nominal–Collaborative	.09	54.49	<0.001	11.25

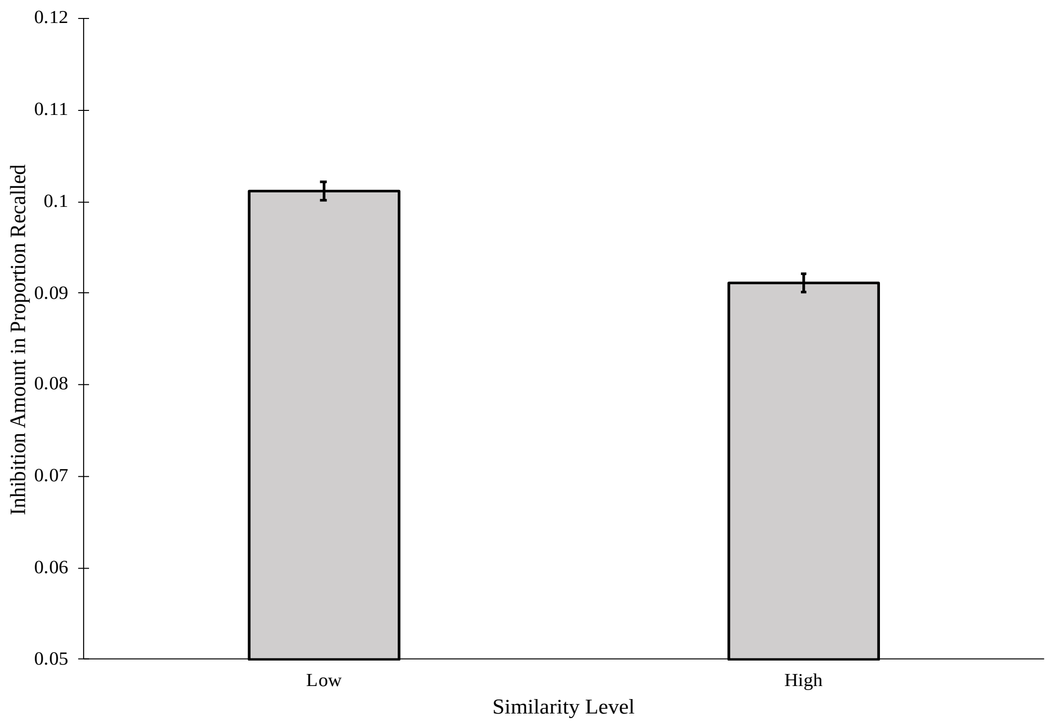

Figure 3.14. *Inhibition amount (mean difference between nominal and collaborative proportion recalled) in low and high background similarity conditions.*

These results show that as background similarity between group members increases, the average proportion recalled in both nominal and collaborative groups decreases. However, when groups have more background knowledge in common, collaborative inhibition is reduced. The results of the ANOVA show that both collaborative groups and nominal groups are disadvantaged by high similarity backgrounds. The interaction effect suggests that memory similarity and homogenization may play a larger role in producing the collaborative inhibition effect than previously thought. These results suggest that shared background knowledge alone can cause a reduction in collaborative inhibition.

2.3 Discussion

The study in this section simulated the effect of shared background knowledge on collaborative inhibition in a controlled manner using the cSAM framework. In the real world, the amount of communication proficiency can vary between different areas of expertise, groups of friends, and married couples. The level of previous communication abilities in all groups is difficult to measure, thus simulating this scenario with cSAM has provided valuable insight into which factors may have contributed to reversing or reducing collaborative inhibition in the original behavioral experiments. The results of this study show that shared background knowledge can reduce collaborative inhibition. These findings support predictions made by the retrieval disruption hypothesis: that shared memory organization provides fewer opportunities for disruption causing a decrease in collaborative inhibition.

These results complement findings in section 1 of this chapter that suggest learning during retrieval decreases the overall proportion recalled by cSAM models. While collaborative inhibition did decrease in the high background similarity condition, the overall proportion recalled by both groups was much lower than in the low background similarity condition. The mechanism responsible for the overall increase in proportion recalled in the high background similarity condition is the same mechanism responsible for the effect found in Figure 3.5. In this figure, collaborative inhibition was still present in the case where learning was prevented during retrieval, however, the overall proportion recalled increased for both collaborative and nominal groups. As model memories become more similar to each other, it is more difficult for models to produce new cues. This is why, when learning was prevented in section 1 of this chapter, the models both produced more responses overall. Similarly, when model memories start off more similar to each other, both nominal and collaborative groups struggle to produce as many responses. In this case, however, collaborative groups with high background similarity gain an advantage over collaborative groups with low background similarity, likely due to a lessened effect of retrieval disruption.

The results of this study support the retrieval disruption hypothesis. When model memories are more similar to each other collaborative models are able to produce more responses and are less affected by incongruous cues from group members. Additionally, learning during retrieval has a negative effect on overall production in both nominal and collaborative groups. By removing the confounding communication factor, these simulations show that shared background knowledge on its own can reduce the effect of collaborative inhibition. It is possible that communication proficiency compounds with shared background knowledge, especially in the case of expert pilots, to produce a faciliatory effect of collaboration. Further efforts to separate communication from shared background knowledge should be pursued to determine the effect of communication proficiency in isolation.

3. The Effect of Group Size

Several studies have shown that as group size increases collaborative inhibition also increases. Originally within the brainstorming literature, Bouchard and Hare (1970) tested the idea output of five, seven, and nine person brainstorming groups, and found the magnitude of the productivity loss increased as group size increased. Historically, collaborative inhibition has consistently been shown to affect groups of three (B. H. Basden et al., 1997; Blumen & Rajaram, 2008; Weldon & Bellinger, 1997) and studies that compared groups of three to groups of four found that larger groups were more affected by collaborative inhibition (B. H. Basden et al., 2000; Thorley & Dewhurst, 2007). These experimental results combined with the findings from the brainstorming literature suggest that as group size increases, collaborative inhibition should increase as well. These results are also predicted by the retrieval disruption hypothesis: as more group members are added, there is a higher chance for individuals to have their retrieval strategies disrupted because more group members are providing external cues.

However, recent studies have shown conflicting evidence for the effect of group size on collaborative inhibition. Luhmann and Rajaram (2015) manipulated group size in their study investigating information transmission in networks using an agent-based model and found that the relationship between collaborative inhibition and group size was not entirely direct. In this study, nominal and collaborative groups of agents were tasked with performing free recall of a list of 40 words in groups of size 2 to 128. They found that as group size increased from 2 to 7, collaborative inhibition increased. But, as group size increased past 7, collaborative inhibition decreased. The authors explain that this effect is caused by nominal groups reaching the recall ceiling much earlier than the collaborative groups. Nevertheless, these findings were predicted and supported by the retrieval disruption hypothesis.

On the other hand, Gates et al. (2022) have found that collaborative inhibition is not present in groups larger than 4. This is the first experimental study that tested groups of sizes larger than 4, and contrary to previous studies, there was not a significant interaction between group type (nominal vs. collaborative) and group size (2, 3, 4, 8, 16) in either turn-taking or free-for-all recall methods. While there are some factors present in this study that may have contributed to the surprising findings (such as presentation method, study location, and study repetitions), this study undeniably raises questions about the predictions that the retrieval disruption hypothesis makes.

The goal of this section is to investigate the effect of group size manipulation on the collaborative inhibition effect produced by cSAM. To do this, cSAM will recreate the simulations presented by Luhmann and Rajaram (2015) and the free-for-all experiment presented by Gates et al. (2022). CSAM is expected to show that as group size increases, collaborative inhibition increases until the recall ceiling is reached. The null results found by Gates et al. (2022) were surprising and given cSAM's previous support of the retrieval disruption hypothesis there is no reason to predict a null effect of collaborative inhibition in these simulations.

3.1 Results

To begin, the group size simulations presented by Luhmann and Rajaram (2015) were recreated using cSAM. Figure 3.10 shows the proportion recalled by both collaborative and nominal groups and the size of the inhibitory effect per group size. The models were set to recall from a list of length 40 with group sizes ranging from 2 to 128. 100 simulations were performed per group size. As group size increased from 2 to 7 collaborative inhibition increased, but, starting with a group size of 8, collaborative inhibition began to decrease. This pattern is almost an exact replica of the pattern found by Luhmann and Rajaram (2015). The decrease in collaboration starting with group size 8, can be explained by the nominal groups reaching the recall ceiling faster than the collaborative groups.

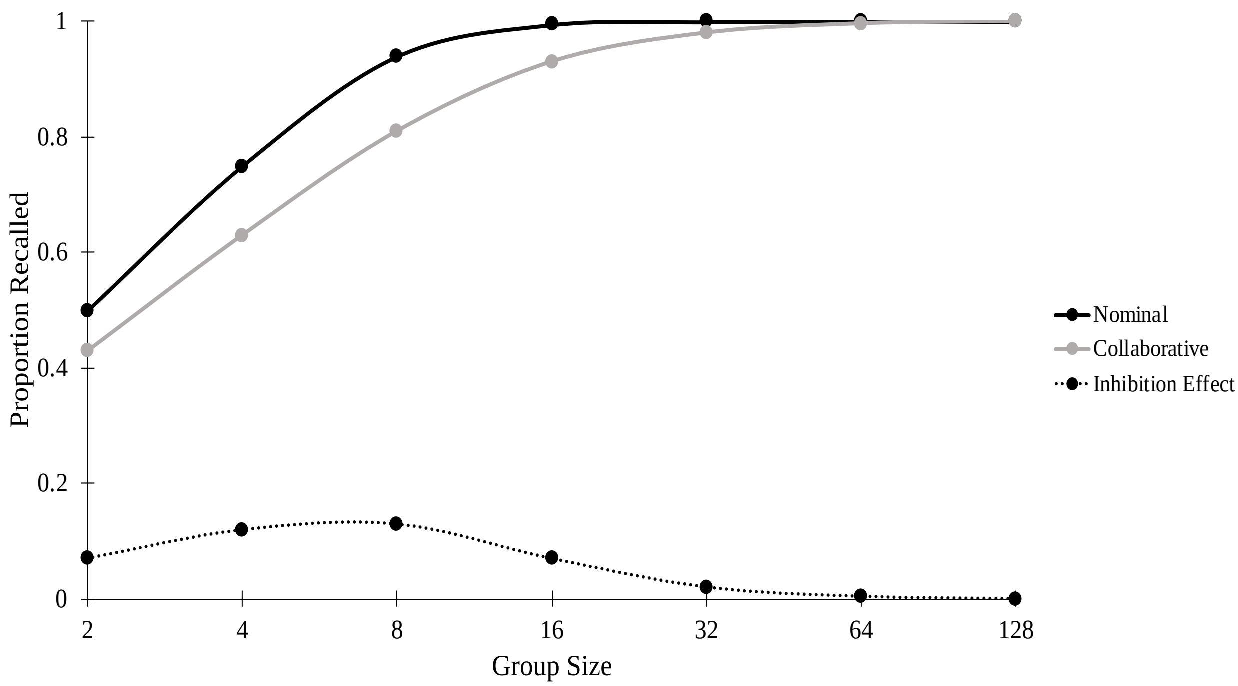

Figure 3.15. *Influence of group sizes 2 to 128 on collaborative inhibition. The x-axis is logarithmic.*

Next, the predictions made by cSAM were compared to the results found by Gates et al. (2022). Figure 3.11 shows the proportion recalled by both collaborative and nominal groups and the size of the inhibitory effect per group size. The models were set to recall from a list of size 60 and in groups of size 2, 3, 4, 8, and 16. 100 simulations were performed per group size. Once again, cSAM showed that as group size increases (up to group size 8), collaborative inhibition also increases. The decrease in collaborative inhibition for group size 16 can be explained by the nominal group reaching the recall ceiling before the collaborative group. Because the list size is longer for this simulation, the nominal groups take longer to reach the recall ceiling.

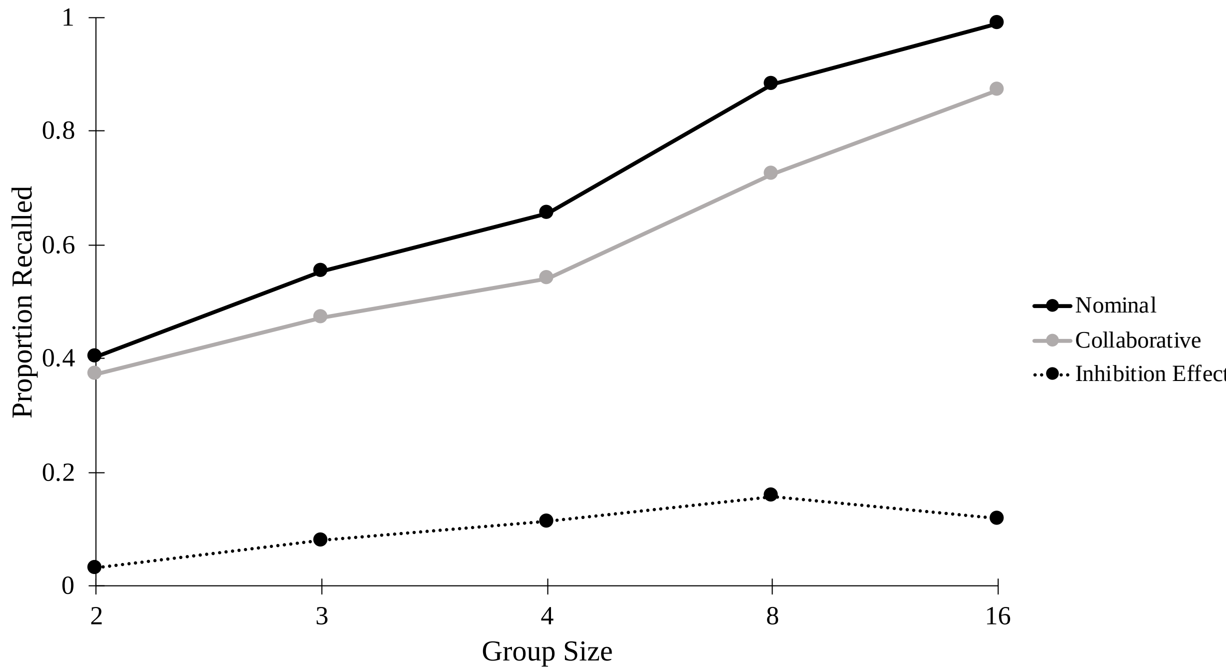

Figure 3.16. *Influence of group sizes 2, 3, 4, 8, and 16 on collaborative inhibition.*

Collaborative inhibition was a significant effect for all group sizes (see Table 3.3). A between-participants two-way ANOVA showed that there were statistically significant main effects of group type, $F(1, 990) = 961.7$, $p < 0.001$, $\eta^2 = 0.49$ and group size, $F(4, 990) = 3671.2$, $p < 0.001$, $\eta^2 = 0.94$, on proportion recalled, both with large effect sizes. Additionally, the interaction between group type and group size was significant, $F\ (4, 990) = 42.9$, $p < 0.001$, $\eta^2 = 0.15$, with a large effect size. These results contradict the findings from Gates et al. (2022) as they only found marginally significant effects of collaborative inhibition in groups of size 3 and 4. Additionally, they found a significant main effect of group size but did not find a significant interaction between group size and group type.

Table 3.5. *Comparison of proportion recalled in nominal and collaborative groups.*

Group Size	Nominal Mean (std)	Collaborative Mean (std)	*df*	*t*	*p*	*d*
2	0.40 (0.05)	0.37 (0.05)	198	4.57	< 0.01	0.60
3	0.55 (0.06)	0.47 (0.05)	198	9.91	< 0.01	1.45
4	0.66 (0.06)	0.54 (0.06)	198	13.4	< 0.01	2.00
8	0.88 (0.04)	0.72 (0.06)	198	22.87	< 0.01	3.14
16	0.99 (0.01)	0.87 (0.05)	198	21.92	< 0.01	3.33

3.2 Discussion

The goal of this section was to investigate which patterns cSAM would show when group size was increased. These results show that as group size increases, so does collaborative inhibition up to the point where groups begin to reach the recall ceiling. These results are predicted by the retrieval disruption hypothesis, which cSAM has shown to support in previous sections of this dissertation. Somewhat surprisingly, cSAM also produced results that were an almost exact replica of the findings presented by Luhmann and Rajaram (2015) even though the two models have a few important differences. First, while their model included psychologically based agents that were able to encode and retrieve memories, the implementation was not as mechanistically extensive as cSAM. Second, the agents performed the collaborative recall task in a turn-taking manner whereas cSAM performs the task using the free-for-all method. Given the results of this study and the finding that memory homogenization plays a large role in producing collaborative inhibition both in cSAM (section 3.1) and Luhmann and Rajaram (2015)'s models, these two modeling approaches seem to make similar predictions.

While cSAM produced results that were predicted by both the retrieval disruption hypothesis and the results presented by Luhmann and Rajaram (2015), they are contradictory to the recent findings of Gates et al. (2022), who found no significant effect of collaborative inhibition on group sizes larger than 4. These findings put the predictions of the retrieval disruption hypothesis under scrutiny, however, there are limitations of the study that may account for the differing results. First, subjects were recruited and tested online, whereas all previous experimental studies that manipulated group size were conducted in person. And second, in the experiment using the free-for-all recall method (which is directly comparable to cSAM), approximately 30% of the data were produced by subjects who had repeated the task. In fact, subjects who participated more than once only improved on subsequent studies when they were in collaborative groups. Meaning, collaborative groups might have performed better in this study than previous studies, causing collaborative inhibition to be less prominent.

Beyond these limitations, the findings from Gates et al. (2022) could suggest some previously unaccounted for mechanistic changes in the collaborative groups. One possibility is that mechanisms other than retrieval disruption are present during collaborative recall and may become more prominent as group size increases. There is some support for this idea in the literature as several mechanisms have been identified to be at work during collaborative recall including retrieval disruption, retrieval inhibition, and memory homogenization (Barber et al., 2015; Luhmann & Rajaram, 2015). An additional reason collaborative inhibition may not be found in larger group sizes, is that larger groups may adopt different retrieval strategies that are less likely to be affected by retrieval disruption. However, additional research will need to be performed to determine whether either of these possibilities are responsible for the surprising findings.

While it is unclear definitively what the relationship between collaborative inhibition and group size is, cSAM predicts that collaborative inhibition should increase as group size increases until group members begin to reach the recall ceiling. This prediction is supported by most previous experimental evidence in addition to the retrieval disruption hypothesis.

4. Post-Collaborative Effects Found in cSAM

The most popular mechanistic hypothesis is the retrieval disruption hypothesis which posits that the inhibitory effects of collaboration occur because individual retrieval strategies are disrupted during group recall (B. H. Basden et al., 1997). However, there is a growing body of experimental evidence that suggests retrieval disruption may not be the only mechanism contributing to collaborative inhibition. According to the retrieval disruption hypothesis, collaborative inhibition should be reduced when group members encode together. Encoding together should lead to more similar memory organization and search strategies, thus reducing the effect of retrieval disruption. Some research has supported this claim (Barber, Rajaram, & Fox, 2012; Finlay et al., 2000; Celia B Harris, Barnier, & Sutton, 2012), but there is a fair amount that has not found this affect (Barber & Rajaram, 2011; Dahlström, Danielsson, Emilsson, & Andersson, 2011), calling into question the role of retrieval disruption in collaborative recall.

Additionally, the retrieval disruption hypothesis predicts that collaborative inhibition should not be present in recognition or cued recall tests because subjects do not rely on their own search strategies when taking these types of tests. Both test types interfere with the idiosyncratic search strategies equally between nominal and collaborative groups, thereby eliminating collaborative inhibition. Again there are conflicting results for this effect, some studies have shown no collaborative inhibition in recognition and cued recall tests (Finlay et al., 2000; Thorley & Dewhurst, 2009), but others have found collaborative inhibition in both recognition (Danielsson et al., 2011) and cued recall tests (M. L. Meade & Roediger, 2009).

Finally, the retrieval disruption hypothesis predicts that when subjects recall non-overlapping lists of words, collaborative inhibition should be reduced. This prediction arises because the cues from other group members don’t directly interfere with the idiosyncratic search strategies of other group members. There have been two studies investigating this effect, the first supported the predictions from the retrieval disruption hypothesis (B. H. Basden et al., 1997) while the second found the opposite effect (Michelle L Meade & Gigone, 2011)—that collaborative inhibition was greater for the unshared list items.

In addition to the growing body of experimental research that suggests other mechanisms may be at play during collaborative recall, cSAM has previously shown evidence for other mechanisms as well: memory homogenization. When subjects recall in groups, their memories naturally become more similar to each other over time. When memories are more similar, it is more difficult to produce new words. This is represented in cSAM by the associative clusters of words becoming more closely related, making it more difficult for the models to “jump” to new associative groups. Given the evidence of alternate mechanistic explanations for collaborative inhibition in cSAM found in section 1 of this chapter, it is essential to consider other possible mechanisms that may affect collaborative inhibition in cSAM.

One such mechanism that is suitable for investigation using cSAM is retrieval inhibition: the tendency for collaborative recall to have a lasting inhibitory effect on group member’s memories. Unlike retrieval disruption, which predicts that there should be a release from the inhibitory effect of collaboration on subsequent individual recall, retrieval inhibition predicts that there should be lasting negative effects from collaboration on subsequent individual recall. Many studies have shown post-collaborative benefits due to re-exposure effects. Subjects performing collaborative recall are exposed to cues they may have forgotten on their own and thus tend to perform better on post-collaborative individual recall (B. H. Basden et al., 2000; Celia B Harris et al., 2012; Weldon & Bellinger, 1997). However, subjects often forget items that they contributed to the previous collaborative recall (Blumen & Rajaram, 2008; Congleton & Rajaram, 2011) indicating lasting negative effects of collaboration. When re-exposure effects were controlled for, the post-collaborative benefits disappeared and only a partial release from collaborative inhibition was found (Barber et al., 2015). The authors interpreted this as support for both retrieval inhibition and disruption because there was only a partial release from inhibition. The goal of this section is to determine whether retrieval inhibition may cause lasting effects of collaboration in cSAM by measuring model performance on post-collaborative individual recall.

4.1 Results

To begin, post-collaborative effects in cSAM were investigated by having groups of models perform collaborative and nominal recall as usual on a list of size 50. Afterwards, the models performed individual recall. Figure 3.12 shows the effects of collaboration on post-collaborative individual recall. Models that previously recalled collaboratively were found to perform significantly better than models that previously recalled in nominal groups, $t(198) = 12.8$, $p < 0.001$. These results are supported by the literature and suggest that cSAM naturally provides post-collaborative benefits because of re-exposure effects.

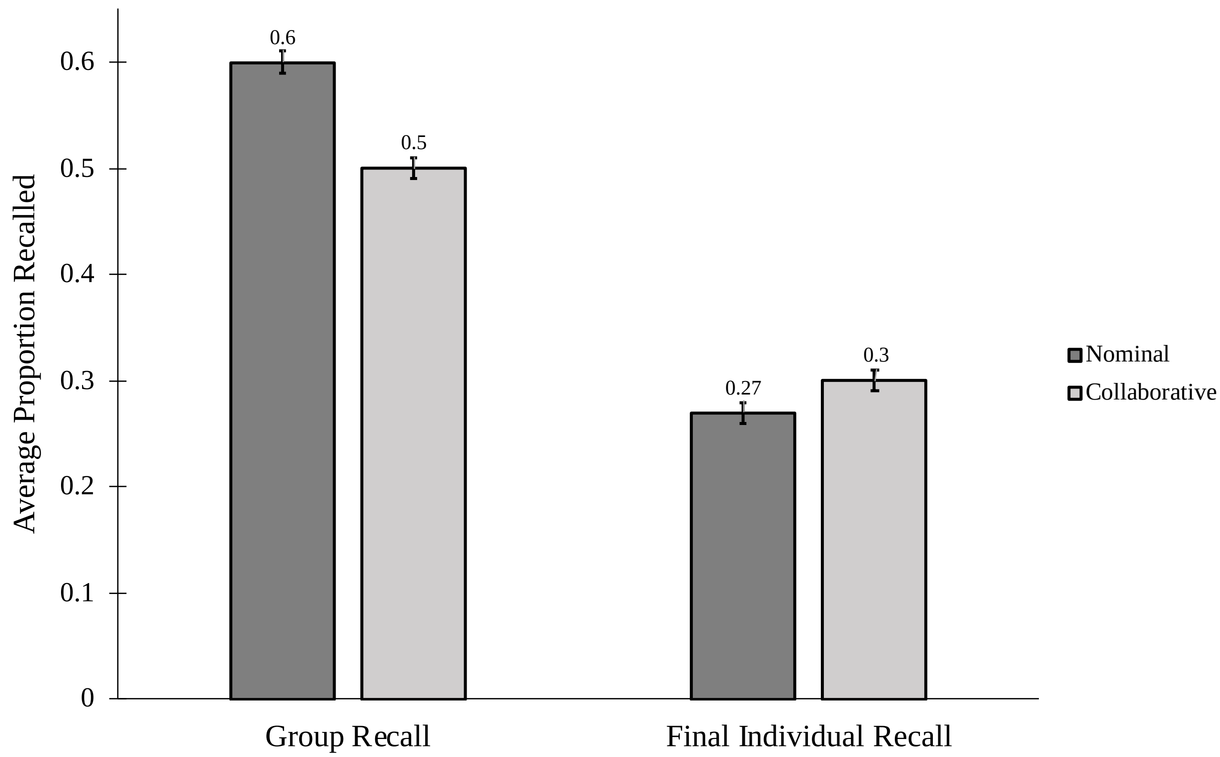

Figure 3.17. *Comparison of collaborative and nominal cSAM groups on a subsequent final individual recall. The collaborative group members experience beneficial post-collaborative effects on individual recall because of re-exposure during collaborative recall.*

To tease apart re-exposure benefits and investigate possible long-lasting deficits due to retrieval inhibition, cSAM was altered so that each group member studied non-overlapping portions of the study list. The models studied non-overlapping portions of one 150-word list (each model only encoded 50 words). This design is a replica of experiment 1 conducted by Barber et al. (2015). Memory deficits present during collaborative recall persisted on subsequent individual recall tests. Table 3.4 compares the proportion correctly recalled by individual group members (out of 50 words) in both nominal and collaborative groups from the experimental study (Barber et al., 2015) and cSAM. The average proportion correctly recalled by the entire group can be calculated by multiplying the average individual proportion recalled by 3 and dividing by 150 (since all group members recalled from non-overlapping lists).

Table 3.6. *Proportion correctly recalled by individual group members in a group recall test and a final individual recall test.*

	Experimental Results		Modeling Results	
Group	**Group Recall (std)**	**Final Individual Recall (std)**	**Group Recall (std)**	**Final Individual Recall (std)**
Nominal Members	0.25 (0.13)	0.25 (0.14)	0.27 (0.03)	0.27 (0.03)
Collaborative Members	0.18 (0.11)	0.20 (0.12)	0.21 (0.03)	0.24 (0.03)

Models that previously recalled in nominal groups recalled a larger proportion of their studied list than models that previously recalled in collaborative groups, $t(198) = 8.96$, $p < 0.001$. A two-way ANOVA showed significant main effects of group type (nominal vs. collaborative), $F(1, 396) = 224.3$, $p < 0.001$, $\eta^2 = 0.362$, and test type (group recall vs. final individual recall), $F(1, 396) = 207.04$, $p < 0.001$, $\eta^2 = 0.085$, with large and medium effect sizes, respectively, on proportion recalled. Additionally, the interaction between the group type and the test type was significant, $F(1, 396) = 9.632$, $p = 0.002$, $\eta^2 = 0.024$, with a small effect size. Figure 3.13 shows the interaction effect between group type and task type, with collaborative groups underperforming compared to nominal groups on post-collaborative individual recall. These results show that when re-exposure benefits are controlled for, collaborative models experience lasting negative effects of collaboration.

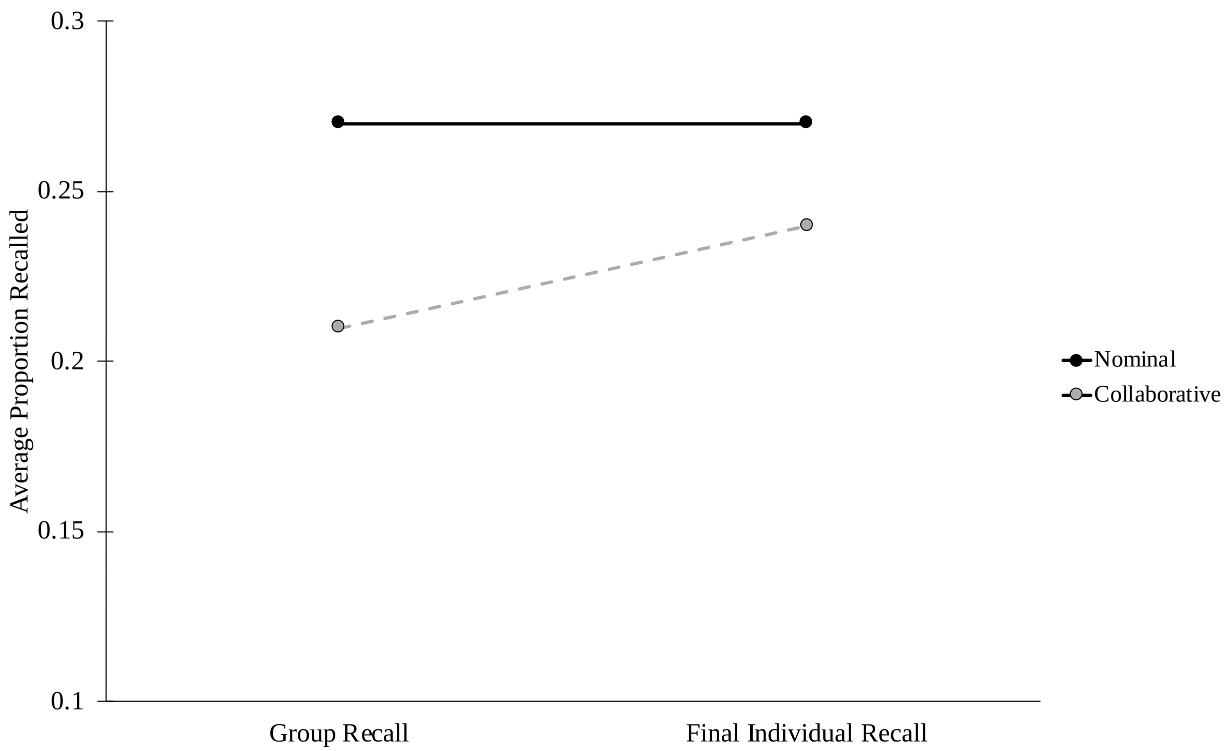

Figure 3.18. *Comparison of collaborative and nominal cSAM group members on a subsequent final individual recall when groups recall from non-overlapping lists. The collaborative group members experience only a partial release from collaborative inhibition in subsequent individual recall.*

4.2 Discussion

The goal of this section was to investigate two known post-collaborative effects of collaborative recall in cSAM. The first effect observed was the post-collaborative benefit of re-exposure effects during collaborative recall. When people recall items from the same study list in groups, they are re-exposed to study items that may have been forgotten by the individual which then boosts individual performance in post-collaborative recall (Celia B Harris et al., 2012). When the cSAM models performed post-collaborative individual recall, the models that previously recalled in collaborative groups recalled significantly more study items than models that previously recalled in nominal groups. Re-exposure benefits in cSAM most likely manifest as increased associations between words which would otherwise be very difficult for an individual model to recall due to low associations with other words. That is, each model starts with unique clusters of associated words, thus increasing associations between previously un-clustered words (via receiving this word as a cue from another model) will make it easier to access in subsequent individual recall. The second effect observed was a partial release from collaborative inhibition when re-exposure benefits were controlled for. When people recall non-overlapping items from a study list, no re-exposure occurs, allowing researchers to tease apart the effects of collaboration from re-exposure. Barber et al. (2015) found that when subjects recalled items from non-overlapping lists, collaborative groups experienced only a partial release from collaborative inhibition on subsequent individual recall. This same pattern was observed in cSAM, showing that collaborative recall has lasting negative effects on individual memory.

A partial release from collaborative inhibition suggests that mechanisms other than retrieval disruption may be active during collaborative recall in cSAM. According to the retrieval disruption hypothesis, collaborative group members should only experience collaborative inhibition during collaboration with no lasting negative effects on memory. However, the retrieval inhibition hypothesis predicts that re-exposure to cues during collaborative recall causes inhibition of other items, which then persists in post-collaborative individual recall. The fact that only a partial release from collaborative inhibition was observed suggests that both retrieval disruption and retrieval inhibition may be active during collaborative recall in cSAM.

In this chapter, three separate mechanisms contributing towards collaborative inhibition have been independently observed in cSAM: retrieval disruption, retrieval inhibition, and memory homogenization. When models recall in a collaborative group, they experience more memory homogenization than models that recalled in nominal groups (see section 1 of this chapter). Increased memory homogenization makes it more difficult for models to produce new responses, thus contributing to collaborative inhibition. In this section, both retrieval disruption and retrieval inhibition have been shown to be present during collaborative recall. While all three of these mechanisms have been observed in cSAM there is experimental evidence that different mechanisms may contribute more heavily towards collaborative inhibition depending on the conditions of the study.

For example, when encoded study items have a high degree of inter-item association, the part-list cuing effect is caused by retrieval disruption. On the other hand, when encoded study items have a low degree of inter-item association the part-list cuing effect is caused by retrieval inhibition (Bäuml & Aslan, 2006). Additionally, retrieval disruption may play less of a role when study items are non-overlapping. Receiving unshared cues during recall may still disrupt idiosyncratic search strategies, but theoretically the disruption should be less than if the cues were from a shared list. The degree to which retrieval disruption and retrieval inhibition affect collaborative inhibition and in which circumstances are largely unstudied and will require further attention to tease apart.

5. Conclusion

The overall goal of this chapter was to investigate several applications of the cSAM framework to gain a deeper understanding of mechanistic causes of collaborative inhibition in cSAM while also making predictions about understudied effects in the experimental literature. First, cSAM was used to investigate the effect of memory homogenization on collaborative inhibition. Collaborative groups experienced increased memory homogenization when compared to nominal groups. Because increased memory homogenization makes it more difficult to produce new items, this may be a contributing factor to collaborative inhibition. However, when learning during recall was prevented and no memory homogenization occurred, collaborative inhibition decreased but did not completely disappear. This suggests that while memory homogenization is not the sole cause of collaborative inhibition it may compound with other cognitive mechanisms like retrieval disruption.

Second, cSAM was used to investigate the effect of shared background knowledge on collaborative inhibition. There is some evidence in the literature that previous relationships and expertise can reduce or reverse collaborative inhibition. However, communication proficiency is heavily entangled with both expertise (M. L. Meade et al., 2009) and previous relationships between participants (Andersson & Ronnberg, 1995, 1996; Johansson et al., 2000). CSAM was used to disentangle communication proficiency and shared background knowledge in section 2 of this chapter. Collaborative inhibition persisted but was reduced when background knowledge was shared, supporting predictions of the retrieval disruption hypothesis. These results suggest that shared background knowledge alone can decrease collaborative inhibition, however, the effect of communication proficiency is still unclear. It is possible that communication proficiency and shared background knowledge combine to decrease collaborative inhibition to an even greater extent, but, this would need further supporting evidence to claim.

Third, cSAM was used to investigate the effects group size on collaborative inhibition. Previously, all experimental and theoretical evidence pointed towards increased group size causing increased collaborative inhibition. CSAM confirmed these predictions and replicated results found in previous modeling attempts (Luhmann & Rajaram, 2015) that were supported by experimental and theoretical evidence. There is one study which contradicts the cSAM findings that shows no increase in collaborative inhibition with the increase in group size (Gates et al., 2022). There are a few concerns about the comparability of that study to previous predictions and the results from cSAM. Beyond these concerns, Gates et al. (2022) suggest that these results may be due to group members adapting different search tactics when in larger groups, or a mechanism other than retrieval disruption may be more prominent in larger groups.

Finally, two post-collaborative effects were studied in cSAM including re-exposure benefits and retrieval inhibition. cSAM was able to naturally show post-collaborative benefits due to re-exposure effects and showed evidence for both retrieval disruption and retrieval inhibition when groups recalled words from non-overlapping lists. The combination of these four applications of the cSAM framework paint a picture that retrieval disruption is not the only mechanism responsible for collaborative inhibition. So far, cSAM has provided supporting evidence for retrieval disruption, retrieval inhibition, and memory homogenization as possible explanations for collaborative inhibition. These results along with other recent studies in the collaborative memory field (Barber et al., 2015; Gates et al., 2022) suggest that more attention should be paid to studying alternative mechanisms to the retrieval disruption hypothesis. The next chapter of this dissertation will detail further applications of cognitive search models similar to cSAM on semantic recall tasks.

CHAPTER 4

COLLABORATIVE VERBAL FLUENCY AND OPTIMAL FORAGING

Collaborative inhibition has been studied almost exclusively in episodic free recall tasks. Early research investigating collaborative inhibition in alternative memory tasks such as recognition memory, cued recall, or semantic memory have been unable to reproduce the effect (Clark et al., 2000; Finlay et al., 2000). Recently, however, collaborative inhibition has been shown in both recognition tasks (Danielsson et al., 2011) and cued recall tasks (Kelley et al., 2012; M. L. Meade & Roediger, 2009). It is unclear why there is a discrepancy between the earlier research and more recent research, however, it brings into question the claim that collaborative inhibition is not present in semantic memory recall tasks.

To date, there is a single peer-reviewed study that has investigated the effect of collaborative inhibition on a semantic memory task. Andersson and Ronnberg (1996) tested groups of friends and non-friends on both episodic and semantic retrieval tasks and found that inhibition was dependent on the memory type. Only explicit, episodic tasks were negatively impacted by collaboration while semantic tasks were not negatively affected. The authors conclude that this result is logical because semantic memory is typically more organized than episodic memory and does not need to be cued with as much precision to retrieve information. Additionally, pilot data for a study was collected by Weldon (2000) which showed a faciliatory effect of collaboration on a general knowledge test. These studies suggest that collaborative inhibition may not be present for semantic recall and that collaboration may be beneficial.

However, both studies measured recall using history questions, some of which were extremely obscure. For example, subjects were asked to answer difficult questions such as, “What American track athlete was the star of the 1936 Olympics?” or, “What is the name of the Chinese religion founded by Lao Tse?” (Weldon, 2000). While these questions certainly probe semantic memory, it is unlikely that subjects have equal knowledge of these history questions—some may not have any ability to answer the history questions at all. Additionally, these types of questions are not necessarily categorically related to each other, which means the group members idiosyncratic organizational structures likely interfere less with recall. Moreover, both studies closely resemble fact retrieval studies which have previously shown that group estimates and event recollection are generally better than individuals, i.e. the “wisdom of crowds” effect (Steyvers, Miller, Hemmer, & Lee, 2009; Surowiecki, 2005). The wisdom of crowds effect is also present when comparing groups which communicate with each other (collaborative) to the aggregate of non-communicative groups (nominal) (Miller & Steyvers, 2011).

Given the similarity of the previous semantic recall studies to fact retrieval studies, it is not surprising that collaboration was beneficial. Thus, a different type of recall task should be used to investigate collaborative inhibition in semantic memory. During a typical episodic free recall task, subjects learn and have access to the same study list and the idiosyncratic search strategies of the group members play a large role in the production of collaborative inhibition. Therefore, a task more comparable to a traditional episodic free recall task is the verbal fluency task. The semantic knowledge probed in the verbal fluency task is relatively simpler, is more likely to be shared equally across subjects, and tends to be highly related to each other with internal organization of categories playing a large role in experimental outcomes (Hills et al., 2012).

Verbal fluency tasks are often used for basic psychological research in addition to clinical tests for early onset dementia (Monsch et al., 1992), Parkinson's Disease (Henry & Crawford, 2004), and schizophrenia (Lundin et al., 2020). Typically, in a verbal fluency task, participants are asked to name as many words as possible in a category within a set time frame. For example, participants could be asked to name as many animals as possible in 3 minutes. Researchers have found that people tend to produce clusters of words during verbal fluency tasks and categorical free recall (Howard, Jing, Addis, & Kahana, 2007; Troyer et al., 1997), suggesting memory may be organized in a patchy manner with information grouped into clusters.

1.1 Optimal foraging theory

Semantic search models often compare search through semantic memory with the foraging behavior of animals where animals need to maximize the resource gain when foraging through a patchy space (Hills et al., 2012). When resources are patchy, foraging animals must decide on where to forage and how long to forage in each patch. Patches are typically defined as local areas with highly clustered resources with the surrounding areas sparse in resources. Additionally, patches vary in quality—some may have more tightly clustered resources and thus provide a higher initial return of gathered resources per unit of time. Given the varying quality of patches, foragers need to decide when to give up searching one local patch and switch to searching for new patches.

Optimal foraging theory describes the behavior of animals foraging for resources. To maximize individual fitness, an animal must adapt foraging strategies that are the most energy efficient. The *marginal value theorem* (Charnov, 1976) is a method of identifying optimal search behavior. The theorem posits that the optimal strategy to maximize resource gain is for individuals to leave a resource patch when the *marginal value* of resource gain within the current patch drops below the long-term average resource gain over the entire space (see Figure 4.1).

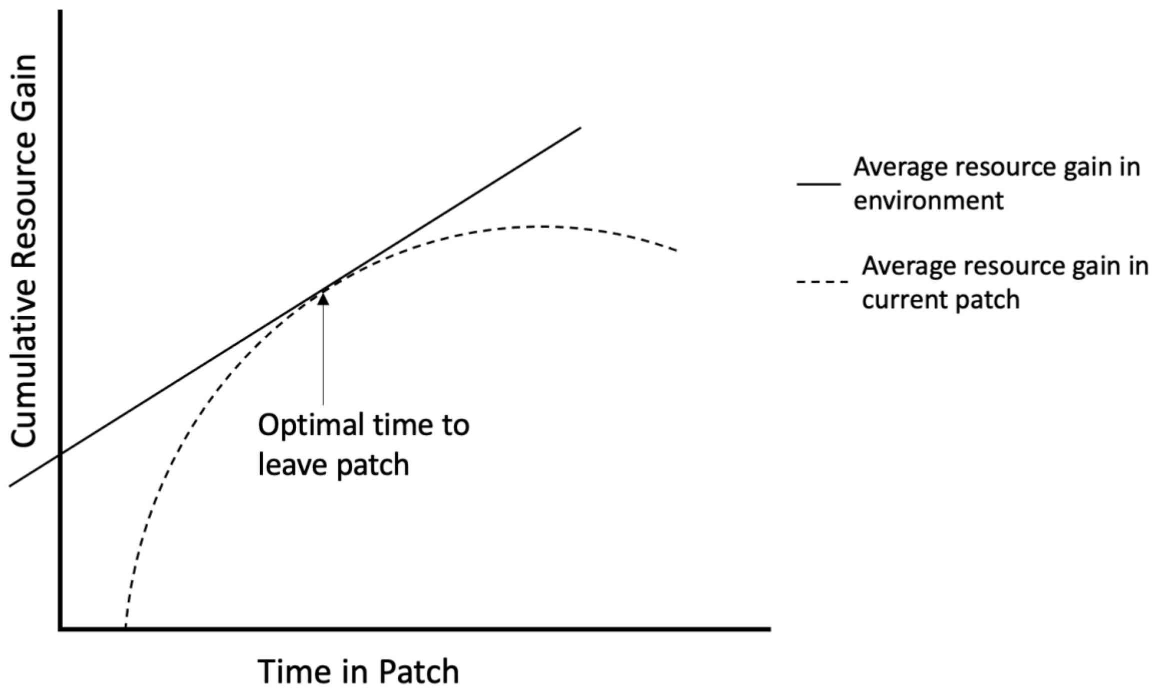

Figure 4.19. *Marginal value theorem. The optimal time to leave a patch is when the line representing resource gain from the current patch creates a tangent with the line representing the average resource gain in the environment.*

For example, when an animal first enters a patch, resource gain will be high because there are abundant resources that are easy to find. However, as the animal continues foraging, resources are depleted and it takes increasingly more effort to find additional resources. According to the marginal value theorem, when the resource rate of return from the current patch drops below the average global rate of return, the animal should stop foraging in the current patch and start looking for a new patch.

Ideas from the optimal foraging field have often been applied to the study of human cognition. Hills (2006) claims that there is an evolutionary link between animals and human foraging behaviors for external resources (such as food) and foraging behaviors for internal resources (such as memory search); that, "What was once foraging in a physical space for tangible resources became, over evolutionary time, foraging in cognitive space for information related to those resources" (Hills, 2006, p. 4). In fact, humans have been shown to forage in a manner consistent with the marginal value theorem when searching for external resources (i.e. visual foraging; Hutchinson, Wilke, & Todd, 2008; Wolfe, 2013) and internal resources (i.e. memory search; Hills et al., 2012; Wilke, Hutchinson, Todd, & Czienskowski, 2009). Additionally, Hills et al. (2012) found that subjects performing a verbal fluency task produced responses reminiscent of foraging for external cues. Thus, they proposed a model of memory search based on the optimal foraging theory, characterizing semantic memory as a cognitive environment with resources distributed in a patchy manner.

1.2 Modeling memory search

There are many parallels between optimal foraging theory and modern models of both episodic and semantic memory search. Models of memory search, including the associative search method used by cSAM, often navigate the patchy nature of memory by modulating between local and global cues. Models begin with a global search to find a cluster of information and then switch to a local search to recall information from that cluster. In cSAM, models begin recall by using a global cue (context-only cues) and then switch to a local cue (word-cue and context-cues) to search within associatively related clusters of study items. When the current cluster resources are depleted (*Lmax* is reached) the model switches to a global cue again (returns to context-only recall) to find a new cluster to repeat the search process. While cSAM applies this search method to episodic recall, similar search models have been applied to semantic memory search.

Verbal fluency tasks are commonly used to observe search patterns within semantic memory. The cluster-switching hypothesis (Troyer et al., 1997) describes fluency lists produced during verbal fluency tasks as an alternating series of clusters (words in a semantic subcategory) and switches (transitions to another subcategory). Searching through a cluster in this case uses a local cue for search while producing a switch uses a global cue for search. This search strategy is almost identical to cSAM's however the global and local cues differ as they are specific semantic memory instead of episodic memory. Originally, verbal fluency data were analyzed using hand-coded category norms indicating which words belonged in which subcategories (Troyer, Moscovitch, Winocur, Leach, & Freedman, 1998). However, the hand-coded norms are often difficult to use in practice as many words produced by subjects are out-of-vocabulary and can exist in multiple subcategories at once. In recent years, cue-switching models have been developed to model search methods within verbal fluency tasks that leverage the increased vocabulary size of distributional semantic models (Hills et al., 2012). The word embeddings of BEAGLE (Jones & Mewhort, 2007) or Word2Vec (Mikolov, Sutskever, Chen, Corrado, & Dean, 2013) are often used in these models to quantify semantic similarity between words and calculate global word frequency within a natural language corpus.

1.3 The forager model

Analysis of verbal fluency data has gone through many changes, from hand-coded norms to integrating vector representations of semantic space produced by distributional semantic models. Most recently, Apsel, Zhang, Kumar, and Jones (2022) have created a Python package, *forager*, specifically for analyzing verbal fluency data. The package is an implementation and expansion of the original foraging model introduced by Hills et al. (2012). The search process used in their model is similar to the cognitive foraging process that modulates between local and global cues which has been previously employed by both the SAM and ACT-R models (Anderson, 2014). Their model assumes that recall works by probing long-term structures in memory with a memory probe, i.e. a set of cues. Additionally, they evaluated two versions of their model, a static and dynamic model. The static model uses the same set of cues during the entire retrieval process which essentially ignores the patchy structure of semantic memory. Meanwhile, the dynamic model makes use of the patchy nature of memory and changes the memory probe depending on whether search is happening locally or globally. When searching within a patch, the dynamic model uses a similarity-based cue (similarity to the previously produced item) and when searching between patches, the model uses a frequency-based cue (context-based search cue). They found that the dynamic model outperformed the static model when modeling verbal fluency data. *Forager* is based on this optimal foraging model and includes options for both static and dynamic model versions.

In addition to a search process, a structural representation of semantic memory is necessary for a model of the verbal fluency task. The original model used both hand coded norms (Troyer et al., 1997) and semantic representations from BEAGLE (Jones & Mewhort, 2007) to construct the search space. *Forager* constructs a semantic search space by using the semantic embeddings from a more modern distributional semantic model, Word2Vec (Mikolov et al., 2013). The animals included in the search space were gathered from words in the *animals* dataset (He, Richie, & Bhatia, 2022; Hills et al., 2012) and any out-of-vocabulary words produced during the verbal fluency task found by Hills et al. (2012), consisting of 1,767 total animal words. A similarity matrix was calculated between all animal words using the semantic vectors from Word2Vec. The frequency values for each word were calculated by querying the Google Books Ngram dataset. The frequency values used in the *forager* model is the logarithm of the raw count for each word in the Ngram dataset.

As well as static and dynamic model options, the *forager* package also includes the option of using various switch methods including: Troyer norms, similarity drop, multimodal similarity drop, and delta similarity. The switch method is a way to identify switches between local and global search. Originally, switches were identified by hand using the Troyer norms, however, Hills et al. (2012) formulated an automatic method of switch identification that they coined the similarity drop model. The similarity drop model predicts a switch if there is a drop in semantic similarity between consecutive items followed by a rise in similarity. This model accounts for differences in clustering structure between participants, but is limited because it cannot produce consecutive switches and may be too sensitive to small similarity variations within clusters (Lundin, 2022). The multimodal similarity drop model is an extension of the similarity drop model which accounts for phonological similarity between words. The similarity between two words is represented as a weighted sum of the semantic and phonological similarities. Finally, the delta similarity switch model is an improvement over the similarity drop model first introduced by Hills et al. (2012). The delta similarity model alternates between clusters and switching if there is a big enough fall in similarity (to go from clustering to switching) or if there is a big enough rise in similarity (to go from switching to clustering). The delta similarity model uses two parameters, rise and fall, as thresholds to designate a word as a switch or a cluster which can be estimated at the participant level. Thus, this method can account for consecutive switches if there is not a big enough rise in similarity to transition from switching to clustering. For the purposes of this study, the dynamic semantic foraging model and the delta similarity switch method will be used.

1.4 The present study

The goal of the present study is two-fold. First, to collect novel data of collaborative groups performing a verbal fluency task and compare performance to nominal groups of the same size. The verbal fluency task is likely more similar to episodic free recall tasks than history quizzes, thus, it is hypothesized that collaborative inhibition will be found in this study. Second, to analyze possible differences in search behavior between nominal and collaborative groups by fitting foraging models to the verbal fluency data. There is a strong connection between foraging models and to cSAM as both make use of local and global cues to model memory search. Given the success of cSAM in modeling collaborative episodic recall tasks, the foraging models are likely to have similar success modeling collaborative semantic tasks. Additionally, the results of this study may provide deeper insight into the connection between cognitive mechanisms responsible for collaborative inhibition in episodic recall tasks and semantic recall tasks (such as retrieval disruption).

2. Methods

2.1 Participants

Participants were 120 undergraduates at Indiana University, Bloomington who received course credit as compensation for participating in this study. Sixty participants were placed into the nominal groups and 60 were placed into collaborative groups. Each group consisted of 3 members, for a total of 20 nominal groups and 20 collaborative groups.

2.2 Procedure

Participants were split into two conditions based on recall method: collaborative and nominal. Subjects were asked to perform both verbal and letter fluency tasks for this experiment. Before beginning the verbal and letter fluency tasks, participants in both collaborative and nominal groups were asked to produce items from a trial category (movie categories) for 1 minute to acclimate to the tasks. Then, participants in both groups were asked to produce items from three semantic categories (animals, foods, vehicles) and three letter categories (words starting with F, A, and S, respectively), and were given three minutes for each task.

The categories were presented one at a time in a randomized order. Participants in the collaborative group were seated in the same room and asked to work together to perform each task. One participant in the group was designated as the scribe and typed the responses produced by the group as they were said aloud on a computer provided by the experimenter. Participants in the nominal groups were asked to perform each task individually and type their own responses as they produced words. The final fluency score was calculated as the number of unique words produced by the collaborative group and the non-overlapping responses from all members in each nominal group were combined to create the final nominal group response.

3. Results

All analysis in this section will be performed on the *animals* category data as the majority of work within the semantic fluency field focuses on this category.

3.1 Collaborative inhibition in semantic memory

Figure 4.2 shows the comparison of the nominal and collaborative performance on the animals category of the verbal fluency task. Nominal groups (N = 20, M = 74.8, STD = 10.7) produced significantly more words than collaborative groups (N = 20, M = 59.5, STD = 6.1), $t(38) = 5.4$, $p < 0.001$. Thus, collaborative inhibition is present in verbal fluency tasks.

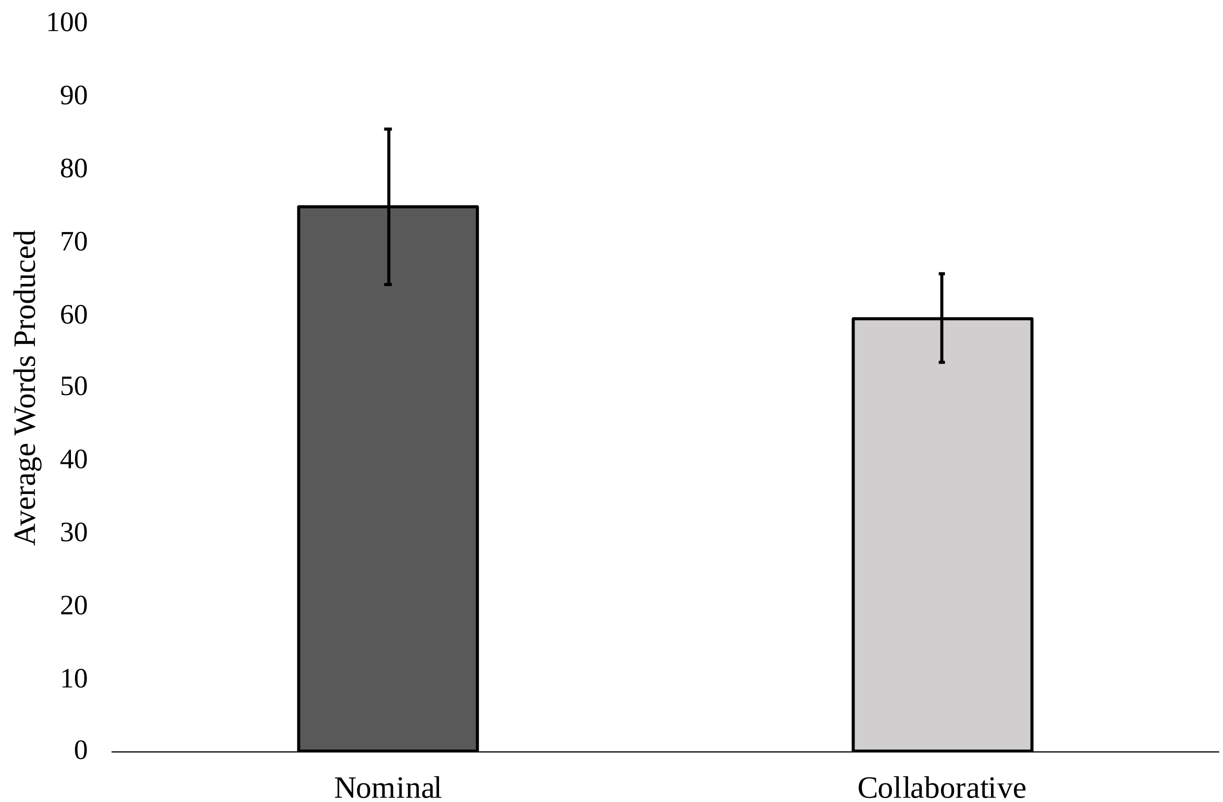

Figure 4.20. *Average words produced by nominal and collaborative groups in the animals category of a verbal fluency task.*

Additionally, the number of repetitions produced by each group were compared. Typically, individuals performing a verbal fluency task will not produce any repeats unless they show signs of dementia (Auriacombe et al., 2006). If collaborative groups repeat words more often, this may indicate participants have more difficulty keeping track of previously produced responses while collaborating. Collaborative groups produced significantly more repetitions (Mdn = 2) than nominal groups (Mdn = 0), $U = 567.0$, $p < 0.001$. The distribution for the nominal individuals was not normal so the medians of the two groups were compared using the Mann-Whitney U test. Because nominal group responses are calculated by combining non-overlapping responses, the nominal groups will naturally have no repeats. To circumvent this and directly compare repeats between groups, repeats by nominal individuals were counted and the median was multiplied by 3 to represent the nominal group. Coincidentally, the median number of repeats for nominal individuals was 0 meaning the nominal groups also had a median of 0 repeats.

3.2 Modeling collaborative search with optimal foraging techniques

In this section, the memory search behavior of nominal individuals, truncated nominal individuals, and collaborative individuals were analyzed. The output from each collaborative individual was recorded during collaborative recall by the experimenter. The truncated nominal individuals were added to the analysis to control for total items produced. To produce the truncated nominal individual data, nominal individuals were paired with collaborative individuals based on total responses provided. So, the nominal individual that produced the most items was paired with the collaborative individual that produced the most items and the nominal individual's list of words was truncated to match the length of the paired collaborative individual. This process was repeated for all nominal and collaborative individual pairs.

The *forager* Python package was used to produce many of the following analyses. Repeated responses were included in this analysis to track subject's search path through semantic space. *Forager* can handle out of vocabulary (oov) words in two ways: by truncating the word list starting at the oov word or by replacing the oov word with the closest match within the vocabulary. To prevent truncating of word lists, oov words without an obvious match in the vocabulary were changed manually to a close alternative that was in the vocabulary. The dynamic model and delta similarity switch method were used to fit this data. Additionally, the rise and fall parameters of the delta similarity switch were estimated at a participant level using the maximum likelihood method.

The mean number of switches (Table 4.1), mean cluster size (Table 4.2), and mean number of clusters (Table 4.3) between groups were analyzed as labeled by the *forager* model. The nominal individuals produced significantly more switches than collaborative individuals, $t(118) = 7.21$, $p < 0.001$ and more clusters than collaborative individuals $t(118) = 6.98$, $p < 0.001$, with no differences in cluster size. These effects disappear when comparing the truncated nominal individuals to collaborative individuals. This suggests that number of items produced is a driving factor for these differences. That is, nominal individuals produce more switches and clusters because they produce more words on average. When items produced is controlled for, there is no significant difference in switching or clustering behavior between nominal and collaborative individuals. Thus, other measures of search strategy will be analyzed to look for more differences.

Table 4.7. *Mean number of switches.*

Group	Mean Switches (std)	Range
Nominal Individuals	12.7 (3.8)	5 – 29
Truncated Nominal Individuals	7.5 (3.2)	2 – 21
Collaborative Individuals	7.9 (3.4)	2 – 15

Table 4.8. *Mean number of clusters.*

Group	Mean Number of Clusters (std)	Range
Nominal Individuals	13.7 (3.8)	6 – 30
Truncated Nominal Individuals	8.5 (3.2)	3 – 22
Collaborative Individuals	9.0 (3.5)	3 – 16

Table 4.9. *Mean cluster size.*

Group	Mean Cluster Size (std)	Range
Nominal Individuals	2.7 (0.7)	1.7 – 5.3
Truncated Nominal Individuals	2.6 (0.7)	1.6 – 5.0
Collaborative Individuals	2.5 (0.8)	1.4 – 5.3

Next, the average pairwise cosine similarity (Table 4.4) of successive other-produced and self-produced words from nominal and collaborative individuals were analyzed. Unlike the previous analyses, the total list length does not need to be controlled for, so truncated nominal individuals were not analyzed. The cosine similarity of each pair of words was calculated based on semantic embeddings from the Word2Vec distributional semantic model (Mikolov et al., 2013). Nominal individuals have no interaction with others during recall, so the other-produced analysis was left blank for both nominal groups. Collaborative individuals trended towards higher pairwise similarity of other-produced items than self-produced items $t(118) = 1.7$, $p = 0.07$. Additionally, collaborative individuals had significantly lower self-produced pairwise similarity when compared nominal individuals, $t(118) = 7.84$, $p < 0.001$. This suggests that collaborative individuals may be paying more attention to items produced by group members than items produced by themselves and that their attention to self-produced items is inhibited compared to nominal individuals.

Table 4.10. *Pairwise cosine similarity of other-produced and self-produced items.*

Group	Other-Produced Similarity (std)	Self-Produced Similarity (std)
Nominal Individuals	--	0.44 (0.03)
Collaborative Individuals	0.41 (0.03)	0.39 (0.04)

The overall frequency (Table 4.5) of other-produced and self-produced words from each group were analyzed. The frequency of each word was also calculated based on semantic embeddings from the Word2Vec model. The frequency of a word is a measure for how common that word is in the Word2Vec vocabulary. There were no differences in frequency between other-produced and self-produced words for the collaborative individuals. However, collaborative individuals produced significantly less frequent, self-produced words than nominal individuals, $t(118) = 3.02$, $p = 0.003$. Collaborative individuals producing less frequent items was not an expected result and suggests that collaborative individuals may be missing more common items or induced to produce less common items by group members.

Table 4.11. *Overall frequency of other-produced and self-produced words.*

Group	Other-Produced Frequency (std)	Self-Produced Frequency (std)
Nominal Individuals	--	6.08 (0.17)
Collaborative Individuals	5.9 (0.15)	5.9 (0.25)

Next, an analysis was conducted to determine whether individuals within collaborative groups were more likely to ignore switches than nominal individuals. That is, one group member attempts to switch, but other group members fail to follow and continue recalling within the current cluster. In the case of nominal individuals, this would present itself as a switch to a new category but then immediately "backtracking" and deciding to go back to the previous cluster. These behaviors would be labeled as multiple switches in a row.

If group members were recalling within a cluster and one group member produced a switch which was ignored by others in the group, then that sequence of 3 words would be labeled as 0 (non-switch, the last word within the original cluster), 1 (the first switch), and 1 (the second switch, either back to the original cluster or to a new cluster). If the group were returning to the first cluster, then the similarity between the last word in the original cluster and the second switch should be higher than the similarity between the last word and the first switch. If, however, the group was not ignoring the switch, then the similarity between the last cluster word and the second switch should not be higher than the similarity between the last cluster word and the first switch. See Figure 4.3 for a visual representation of this scenario.

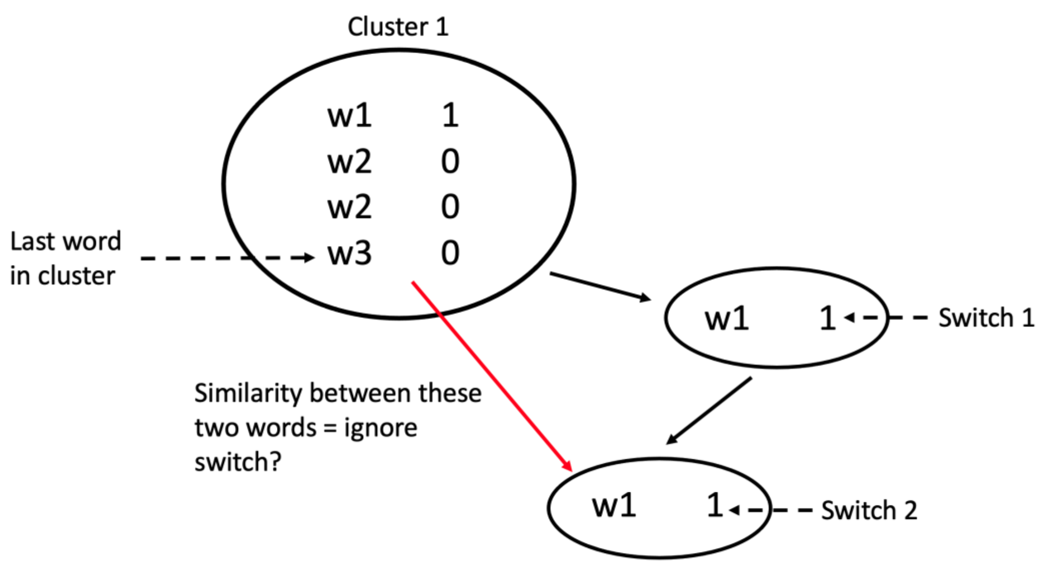

Figure 4.21. *Switching behavior of groups producing multiple switches in a row.*

To determine whether multiple switches in a row indicates switches being ignored or is simply a fast switch between semantically distinct clusters, the average similarity between the last item in a cluster and the second switch was compared to the average similarity between the last item in a cluster and the first switch for every instance where 2 switches occurred sequentially. The switching behavior of collaborative groups as a whole was included in this analysis as group members may be more likely to pay attention to the words produced by other group members. Analyzing the collaborative individuals alone would ignore relevant context provided by the group as a whole. Truncated nominal individuals were also included in this analysis, as switching and clustering behavior is expected to change when more items are produced.

Table 4.12. *Mean similarity between last word in a cluster and the 1st and 2nd switches and mean occurrence of multiple switches.*

Group	Mean similarity last word in cluster— 1st switch (std)	Mean similarity last word in cluster— 2nd switch (std)	Mean number of multiple switches (std)
Nominal Individuals	0.35 (0.04)	0.37 (0.07)	5.1 (3.3)
Truncated Nominal Individuals	0.36 (0.07)	0.39 (0.12)	2.8 (2.2)
Collaborative Individuals	0.32 (0.08)	0.35 (0.10)	3.3 (2.5)
Collaborative Groups	0.35 (0.04)	0.41 (0.05)	10.1 (5.3)

Collaborative groups, $t(19) = 4.48$, $p < 0.001$, and truncated nominal individuals, $t(53) = 2.46$, $p = 0.02$, had significantly higher mean similarity between the last word in a cluster and the second switch than the last word in a cluster and the first switch. However, there was not a significant difference in the similarity of the last word in a cluster and the second switch between truncated nominal individuals and collaborative groups. These results suggest both collaborative groups and nominal individuals have some degree of "ignoring" switches going on. But, collaborative individuals don't seem to do this more often or to a higher degree than truncated nominal individuals. Additionally, nominal individuals had significantly more instances of multiple switches than collaborative individuals, $t(109) = 2.20$, $p = 0.03$, however, this effect disappears when controlling for number of words produced with the truncated nominal individuals. This is an expected result as individuals performing verbal fluency tasks tend to switch faster near the end of the task as they run out of easily accessible common word clusters.

It is possible that both groups display this "ignoring switches" behavior because of natural overlap in subcategories. For example, a group could be naming birds and produce the following string of words, "eagle", "hawk", "penguin", "falcon". All of these are birds, however, penguin is less similar to hawk than hawk is to falcon, so the sequence would likely be labeled as 1, 0, 1, 1 with two switches in a row at the end. This is a case where it's difficult to determine category boundaries. It seems reasonable to have 2 switches in a row because "penguin" is not a bird of prey like the others, but on the other hand all 4 words can be classified under one category—birds. It's possible that both nominal and collaborative groups create poorly defined categories that give the appearance of ignoring switches when in reality, the first switch could also be classified as a non-switch given less strict category boundaries.

Finally, the *forager* model output provides an index of relative saliency for frequency and similarity information captured by the beta parameter value. The frequency beta parameter is an indicator of the amount of attention paid to the frequency of words when switching between clusters. Thus, higher values of the frequency beta parameter indicate more attention paid to overall frequency of words when switching between clusters. The similarity beta parameter is an indicator of the similarity between words within a cluster. So, higher similarity beta values indicate the words within clusters are more semantically related. Table 4.7 shows the mean and standard deviation of the beta values for both frequency and similarity by group. Truncated nominal individuals were not included in this analysis as the frequency and similarity betas do not depend on the amount of items produced. Nominal individuals were found to have significantly larger beta values for frequency cues, $t(118) = 2.27$, $p = 0.03$, and similarity cues, $t(118) = 2.42$, $p = 0.02$, than collaborative individuals. However, when comparing nominal individuals to the entire collaborative group, only the frequency beta value was significantly larger in the nominal individuals, $t(78) = 3.13$, $p = 0.002$.

Table 4.13. *Mean frequency and similarity beta values.*

Group	Frequency Beta Mean (std)	Similarity Beta Mean (std)
Nominal Individuals	8.1 (1.2)	4.4 (0.5)
Collaborative Individuals	7.6 (1.5)	4.1 (0.7)
Collaborative Groups	7.3 (0.9)	4.2 (0.5)

The differences in frequency beta values in this analysis are complimentary to the analysis of overall frequency of produced words shown in Table 4.5. Collaborative individuals and groups seem to be producing less frequent words on average and are not paying as much attention to frequency when switching between clusters. Additionally, the differences in the similarity beta values show a similar pattern to the pairwise similarity analysis presented in Table 4.4. Collaborative individuals had lower pairwise similarity overall than both nominal individuals and trend towards having lower similarity beta values than nominal individuals. However, when analyzing the output from collaborative groups as a whole, the similarity beta values were not significantly lower than nominal individuals. This suggests that the collaborative groups are producing more tightly clustered words than collaborative individuals and supports the findings in

Table 4 that collaborative individuals trend towards a higher pairwise similarity between other-produced items than self-produced items.

4. Discussion

The results of this study show that there are several key differences between collaborative and nominal groups in semantic search. First, collaborative groups produce significantly fewer responses in collaborative verbal fluency task than nominal groups. Thus, collaborative inhibition is present in verbal fluency tasks. Second, when the number of words produced was controlled for, there were no significant differences in number of switches, cluster size, or number of clusters between nominal and collaborative individuals. This prompted additional analyses on the differences in search strategies between the two groups. An analysis of average number of repeats revealed that collaborative groups produce more repeats than nominal groups, suggesting collaborative group members have more difficulty keeping track of items produced by group members. An analysis of pairwise similarity and overall frequency of words produced revealed that collaborative individuals tend to pay more attention to words produced by group members compared to words produced by themselves and produced less frequent words on average than nominal individuals. An analysis of whether collaborative groups are more likely to ignore switches and continue clustering revealed that both collaborative groups and nominal individuals tended to produce multiple switches in a row—indicating that this effect may be produced by unclear category boundaries instead of ignoring switches. Finally, an analysis of the frequency and similarity beta values estimated by the *forager* model revealed that collaborative individuals and groups paid less attention to the frequency of a word while switching between clusters and that collaborative individuals tended to produce less tightly related clusters of words.

While collaborative inhibition was predicted to appear in verbal fluency tasks, this result is somewhat surprising considering previous findings. This is the first time collaborative inhibition has been observed in a semantic recall task. Previous studies, of which there were only two, tested general history knowledge and concluded that either no collaborative inhibition was present or collaboration was faciliatory (Andersson & Ronnberg, 1996; Weldon, 2000). However, while history questions do indeed probe semantic memory, it's highly likely that the underlying category structure of such history questions is either not strong enough or does not differ between group members. Retrieval disruption relies on differences in memory structure between group members to produce the inhibitory effect. If there is no strong category organization, which is likely the case when answering obscure history questions, then there is no retrieval organization to disrupt which would lead to no collaborative inhibition. In episodic tasks, when subjects have very similar memory organizations collaborative inhibition is reduced or eliminated (B. H. Basden et al., 1997; Finlay et al., 2000). It follows that the same explanation could apply to recall of history facts. Individuals performing verbal fluency tasks on the other hand, have been shown to produce words with a distinct organizational pattern within each category (Hills et al., 2012). It is possible that the underlying semantic organization of verbal fluency categories could be very similar across subjects as knowledge of semantic categories is more widespread. If this were the case, then collaborative inhibition may not have appeared in verbal fluency tasks. However, the findings of this study suggest that there is sufficient non-overlapping internal organization of semantic categories between subjects to produce collaborative inhibition.

Nominal groups were found to produce fewer repeats on average than collaborative groups. In the past, it has been found that individuals do not produce repeats during a verbal fluency task unless they show early signs of dementia (Auriacombe et al., 2006). Repeating words in the verbal fluency task indicates suboptimal performance from collaborative groups. These repeats could be an indication that group members have difficulty keeping track of what the group has already produced. However, during the task, all group members have access to a list of all words previously produced by the group. So, theoretically, group members should be able to look at the list during the study and filter repeats. Another possible explanation for the increased repeats is that words previously recalled by the group have had their associations increased, causing group members to be more likely to recall them in the future, a mechanism called retrieval blocking. Retrieval blocking is one of the possible mechanisms responsible for the part-list cuing effect and seems like it may be playing a role in collaborative verbal fluency production. Retrieval blocking has not received much attention within collaborative memory research, and the single study that investigated its effects concluded that it was not a driving factor of collaborative inhibition (Barber et al., 2015). Given the possible differences between semantic recall and episodic recall, it is possible that retrieval blocking could play a role in collaborative inhibition within semantic recall tasks. However, further experimental research would need to be conducted to explore this possibility.

The additional analyses performed to further investigate search strategies indicated that collaborative individuals tend to employ suboptimal search strategies compared to nominal individuals, such as producing less frequent words, paying less attention to word frequency, and producing less semantically related clusters than nominal individuals. These findings can be explained by the retrieval disruption hypothesis. If collaborative individuals have their retrieval organizations disrupted by other group members, then they could be induced to remain within clusters for longer, suboptimal periods of time which would decrease their overall output. For example, if a collaborative group is recalling birds and group member 1 has reached the limit of their easily accessible knowledge, then according to the marginal value theorem, group member 1 should switch and find a new category. However, if group member 2 is comparatively more knowledgeable about birds, then they may keep producing bird words that are increasingly less common. Group member 1 may have ideally wanted to switch categories, but is now receiving obscure bird words that force their attention back to the bird category inducing them to continue searching for bird words when they should have switched categories. In this case, group member 1 would be much slower at searching through remaining birds and would produce fewer words than if they had switched categories. Further analysis of reaction times is needed to fully support this possibility. If more time is taken between responses within clusters, then that would provide more evidence supporting this explanation. Collaborative individuals producing less semantically similar words can also be explained by the retrieval disruption hypothesis. Theoretically, each group member has their own internal structure of the animals category with certain words having higher internal associations than others. When provided with suboptimal cues from other group members, they are unable to retrieve the most semantically similar words from their own semantic clusters and are pulled towards somewhat more similar words to the current cue, thus causing collaborative individuals to produce less semantically similar words.

This study has shown that (1) collaborative inhibition is present in collaborative verbal fluency tasks, (2) collaborative individuals employ suboptimal search strategies when compared to nominal individuals, and (3) suggests the presence of both retrieval disruption and retrieval blocking as explanatory mechanisms for collaborative inhibition in semantic tasks. These are novel findings that support the idea that memory structure is extremely important when studying collaborative inhibition. While previous studies investigating collaboration in semantic memory tasks did not find evidence of inhibition, the underlying memory structure for the types of questions asked was likely not strongly organized. However, participants performing verbal fluency tasks have previously been shown to produce outputs in a manner that suggests a strong underlying memory structure of each semantic category. Thus, the verbal fluency task is more appropriate to employ when studying the effects of collaboration on semantic memory. Additionally, other mechanisms, such as retrieval inhibition or memory homogenization are possible contributing factors to collaborative inhibition in semantic tasks. Retrieval disruption, memory homogenization, and retrieval inhibition all have supporting evidence both in experimental studies and in previous chapters of this dissertation for episodic recall tasks. Thus, it is possible that these factors also play a role in collaborative semantic tasks. A cognitive model of collaborative verbal fluency would be greatly helpful in investigating the role of each cognitive mechanism in producing the inhibitory effect in verbal fluency tasks.

CHAPTER 5

CONCLUSION

1. Overview of Dissertation

The goal of this dissertation was to expand the modeling efforts for collaborative memory within both episodic and semantic memory tasks by developing a formal computational model of collaborative recall. This work shows that SAM, adapted to cSAM, can act as a unified theory to explain both individual and collaborative memory effects and can also provide insight into mechanistic explanations of collaborative inhibition that would be difficult or nearly impossible to study behaviorally.

Chapter 1 provided an overview of the collaborative memory field, collaborative inhibition, mechanistic hypotheses of collaborative inhibition and other modulating factors. Additionally, this chapter motivated the creation of a formal modeling framework with which to study collaborative memory effects in both episodic and semantic memory tasks. This chapter also introduced the idea of a multi-process mechanistic account of collaborative inhibition and provides the framing for the rest of the dissertation.

Chapter 2 presented a framework to scale the Search of Associative Memory model (SAM; Raaijmakers & Shiffrin, 1981) to collaborative free recall paradigms with multiple models working together (coined cSAM). SAM was chosen as the base model for the collaborative framework because it is well-studied, is the most widely used model in episodic memory research (Wilson & Criss, 2017; Wilson et al., 2020), and is one of the only cognitive models that has been shown to successfully model the part-list cuing effect in individual memory (Raaijmakers & Shiffrin, 1981). Importantly, each SAM model in isolation still retains the explanatory power for the range of behavioral phenomena in individual memory paradigms, providing a unified model to understand both individual and collaborative memory. cSAM was able to reproduce the basic patterns of collaborative inhibition seen in experimental data including effects found for both categorized and uncategorized lists. When fitting learning parameters to individual data, no significant differences were found between nominal and collaborative groups, indicating that collaborative inhibition in cSAM is caused by a mechanistic or structural difference that is not captured by parameter differences. The simulations in this chapter suggest that collaborative inhibition arises from a mismatch of the cues used to search memory and the differing subjective organizations of the group members, as predicted by the retrieval disruption hypothesis.

While the initial findings in Chapter 2 suggest that retrieval disruption may cause collaborative inhibition in cSAM, recently, the possibility of a multi-process account of collaborative inhibition has gained traction in the experimental literature. **Chapter 3** investigates memory homogenization, shared background knowledge, group size, and post-collaborative effects as they have all been shown to moderate collaborative inhibition (Barber et al., 2015; Luhmann & Rajaram, 2015; M. L. Meade et al., 2009; Thorley & Dewhurst, 2007). Thus, the goal of Chapter 3 was to investigate these areas to gain a deeper understanding of what is driving collaborative inhibition in cSAM while also making predictions about understudied effects in the experimental literature.

The first effect studied in Chapter 3 was memory homogenization: the idea that learning during a collaborative recall task decreases the diversity of group member's memory representations, which then reduces collaborative recall performance and can cause collaborative inhibition (Luhmann & Rajaram, 2015). This effect is not directly observable in behavioral studies as it requires direct access to memory values. cSAM model memories were found to naturally homogenize during recall and when this was prevented, collaborative inhibition persisted, though the effect size was diminished. This suggests that memory homogenization may be a contributing factor to collaborative inhibition along with retrieval disruption.

The second effect studied in Chapter 3 was shared background knowledge. Studies that previously investigated the effect of shared background knowledge on collaboration were unable to differentiate between the effect of shared background knowledge and high levels of communication proficiency. By simulating shared background knowledge in cSAM, the effect of communication proficiency was entirely removed. Results of this section show that shared background knowledge can reduce collaborative inhibition. These findings support predictions made by the retrieval disruption hypothesis: that shared memory organization provides fewer opportunities for disruption causing a decrease in collaborative inhibition.

The third section of Chapter 3 investigated the effect of increased group size on collaborative inhibition. Most previous studies predict that as group size increases, collaborative inhibition increases (B. H. Basden et al., 2000; Luhmann & Rajaram, 2015; Thorley & Dewhurst, 2007), which is a prediction consistent with the retrieval disruption hypothesis. However, a recent study predicted the opposite, that group sizes larger than 4 do not produce collaborative inhibition

(Gates et al., 2022). Given the recent controversy surrounding the effect of group size on collaborative inhibition, investigating this effect with cSAM was worthwhile. Simulations with cSAM showed that collaborative inhibition does increase as group size increases until group members begin to reach the recall ceiling. This prediction is supported by most previous experimental evidence and the retrieval disruption hypothesis.

The fourth and final section of Chapter 3 investigated post-collaborative effects in cSAM such as retrieval inhibition and re-exposure effects. Retrieval inhibition is the tendency for collaborative recall to have a lasting inhibitory effect on group member's memories. Unlike retrieval disruption, which predicts that there should be a release from the inhibitory effect of collaboration on subsequent individual recall, retrieval inhibition predicts that there should be permanent negative effects from collaboration on subsequent individual recall tasks. However, this effect is often offset by the beneficial effects of re-exposure during collaborative recall. cSAM was shown to produce beneficial re-exposure effects for the collaborative models on subsequent individual recall. Additionally, when re-exposure benefits were controlled for by having models learn and recall from non-overlapping portions of a list, cSAM produced only a partial release from collaborative inhibition on subsequent individual recall tasks. Observing a partial release from collaborative inhibition suggests that both retrieval disruption and retrieval inhibition may be contributing to collaborative inhibition in cSAM.

The findings of Chapter 3 further validate cSAM as a model of collaborative recall as it was able to reproduce and provide predictions in line with the experimental literature for all effects studied. Additionally, this chapter further supports a multi-process account of collaborative inhibition as retrieval disruption, retrieval inhibition, and memory homogenization were each found to play a role in producing collaborative inhibition in cSAM.

Collaborative inhibition has been studied almost exclusively in episodic free recall tasks. Studies investigating collaborative inhibition in semantic recall tasks used paradigms more similar to fact retrieval studies which are arguably too dissimilar to compare to episodic free recall tasks. In **Chapter 4**, verbal fluency tasks were adapted to a collaborative setting because the task is a more appropriate comparison to collaborative episodic tasks. Optimal foraging models, which have previously been used to model search processes in verbal fluency tasks, were used to analyze behavioral patterns found in this study. Optimal foraging models use the same search process as cSAM and thus is a good adaptation of the modeling framework for use on a verbal fluency task.

The results of this study show that collaborative inhibition is present in collaborative verbal fluency tasks. This finding is completely novel as collaborative inhibition has never been found before in collaborative semantic memory tasks. The analysis of search behavior using the optimal foraging models revealed several key differences between nominal and collaborative groups. First, collaborative groups produced more repeats than nominal groups, suggesting collaborative group members have more difficulty keeping track of items produced by others. A possible explanation for this result is that words previously recalled by the group have had their associations increased, causing group members to be more likely to recall them in the future, a mechanism called retrieval blocking. Retrieval blocking is not a commonly studied mechanistic cause of collaborative inhibition, however, it may be more likely to play a role in producing collaborative inhibition on semantic memory tasks. Additional analyses indicate that collaborative individuals tend to employ suboptimal search strategies compared to nominal individuals, such as producing less frequent words, paying less attention to word frequency, and producing less semantically related clusters when compared to nominal individuals. These findings are predicted by the retrieval disruption hypothesis.

2. Future Directions

This dissertation has expanded on the formal modeling efforts within the field of collaborative memory. While many topics have been explored in this work, there are additional applications of a formal model of collaborative recall. One such application that is particularly fit for study using formal computational models similar to cSAM, is the creation and spread of false memories in collaborative groups and larger networks. Because collaborative group interactions at smaller scales (such as interactions studied within the collaborative memory field) are believed to create the groundwork for larger scale collective memories (Choi et al., 2017; Maswood & Rajaram, 2019), expanding cSAM to larger networks is a logical next step.

There are two prominent manipulations suitable for studying the creation and spread of false memories using an adapted cSAM model: method of collaboration and group structure. Recent research has shown that false recall in groups may be affected by the method of collaboration. Typically, the method of collaboration in a collaborative memory task is either a turn-taking or free-for-all method. The turn-taking collaboration method requires group members to wait their turn before producing a response and continue recalling until a set amount of time has passed or no group members can produce more responses. The free-for-all collaboration method, on the other hand, allows group members to freely interact with others, produce responses at any given time during the experiment, and correct errors as they see fit. When these two collaboration methods were directly compared using DRM word lists, Thorley and Dewhurst (2007) found that the turn-taking method was more likely to cause false memories in collaborative groups whereas the free-for-all collaborative groups did not share this downfall. This is believed to be because the turn-taking method provides little to no opportunity for group members to correct errors during recall, thus false memories are more likely to emerge under these conditions. The original cSAM framework adapted a free-for-all collaboration method precisely because the turn-taking method has shown to produce more memory intrusions (M. L. Meade & Roediger, 2009; Rajaram & Pereira-Pasarin, 2010). However, if the goal of the research is to investigate the creation of false memories, then adapting cSAM to use a turn-taking method of collaboration may be appropriate.

The structure of a collaborative group has also been shown to have significant effects on the spread of false memories and collaborative inhibition in post-collaborative recall. Choi et al. (2014) investigated the role of memory transmission in collaborative groups with varying group structures. Three conditions with alternate group structures were tested, each condition consisted of 3 rounds of recall. The control group only recalled individually, and the two collaborative conditions had subjects recall in collaborative groups twice before completing a final individual recall. Of the two collaborative conditions, one condition has an insular group structure and the other has a diverse group structure. The insular collaborative condition had subjects recall with the same group twice before performing a final individual recall. The diverse collaborative condition had subjects recall with two completely different groups before performing a final individual recall. The authors found that collaborative inhibition was eliminated in the diverse collaborative condition, likely due to the positive effects of re-exposure outweighing the negative effects of retrieval disruption. Additionally, they found that insular and diverse collaborative groups were equally able to produce shared memories, but diverse collaborative groups lead to less forgetting by the group. These results show that the group structure has a profound effect on both collaborative inhibition and the likelihood of group members misremembering studied items.

Another study using the same paradigm of diverse vs. insular collaborative groups (Choi et al., 2017), found that diverse collaborative groups tended to produce more false memories after collaboration, however these false memories were ranked with low confidence. When testing the collective memory of the group afterwards, not many of the false memories persisted. On the other hand, insular collaborative groups showed increased confidence in false memories and increased presence of false memories in the group's resulting collective memory following collaboration. This suggests that while false memories are present in both insular and diverse group structures, diverse group structures prevent these memories from becoming permanent fixtures in the group's resulting collective memory. Adapting cSAM to investigate differences in group structure could provide valuable insight into the mechanisms responsible for these differences that aren't easily investigated in behavioral studies. For example, memory homogenization was found to be a contributing factor for collaborative inhibition and likely plays a large role in the formation of collective group memories following multiple instances of collaboration.

Finally, this dissertation was the first observation of collaborative inhibition in in a semantic memory task. To date, there has been no experimental research investigating collaborative inhibition in verbal fluency tasks beyond the present work. Given the abundance of possible mechanistic explanations for collaborative inhibition studied for episodic memory tasks, it is reasonable to investigate such explanations for semantic memory tasks as well. However, it is entirely possible that the mechanisms responsible for collaborative inhibition may differ between episodic and semantic memory tasks due to the differences in memory structure. For example, this dissertation, as well as several other studies, have found that retrieval inhibition may play an active role in producing collaborative inhibition (Barber et al., 2015) and the part-list cuing effect (Aslan, Bäuml, & Grundgeiger, 2007; Lehmer & Bäuml, 2018). While analysis of experimental data collected in Chapter 4 pointed towards retrieval disruption and possibly retrieval blocking as mechanisms responsible for collaborative inhibition, retrieval inhibition should not be discounted given its prevalence in episodic free recall tasks. Though, it is likely that retrieval inhibition may not be as prominent in semantic memory tasks as subjects performing the verbal fluency task are not asked to learn new lists of words. It would be surprising if collaboration on a verbal fluency task were to permanently inhibit the memory of various animal words. What is more likely, is that the memory of the words themselves are not inhibited, but that the subjective organization of the animals category is re-arranged to better fit the order produced by the group. This idea has been supported in earlier experimental work where subjects were more likely to reproduce the order of words produced during collaborative recall during subsequent individual recalls (Congleton & Rajaram, 2014).

3. Synthesis of Findings

The cSAM framework has repeatedly been able to reproduce collaborative inhibition patterns found in the experimental literature that support a multi-process account of collaborative inhibition. cSAM has shown evidence for three different mechanistic explanations for collaborative inhibition: retrieval disruption, retrieval inhibition, and memory homogenization. Additionally, results found during the collaborative verbal fluency study in Chapter 4 suggest that retrieval blocking may be a contributing factor to collaborative inhibition in semantic memory tasks. The culmination of these findings suggests that while most experimental researchers have focused on retrieval disruption as the primary cause of collaborative inhibition, future researchers should begin investing time in other mechanistic explanations. Finally, the collaborative modeling framework presented in this dissertation could be expanded to further investigate collaborative inhibition in semantic memory tasks and scaled up to investigate the creation and spread of false memories in collaborative groups and networks for a more complete account of collaborative memory.

References

Anderson, J. R. (2014). *Rules of the mind*: Psychology Press.

Andersson, J. (2001). Net effect of memory collaboration: How is collaboration affected by factors such as friendship, gender and age? *Scandinavian Journal of Psychology, 42*, 367-375.

Andersson, J., Hitch, G., & Meudell, P. (2006). Effects of the timing and identity of retrieval cues in individual recall: an attempt to mimic cross-cueing in collaborative recall. *Memory, 14*(1), 94-103. doi:10.1080/09658210444000557

Andersson, J., & Ronnberg, J. (1995). Recall suffers from collaboration: Joint recall effects of friendship and task complexity. *Applied Cognitive Psychology, 9*, 199-211.

Andersson, J., & Ronnberg, J. (1996). Collaboration and memory: Effects of dyadic retrieval on different memory tasks. *Applied Cognitive Psychology, 10*, 171-181.

Apsel, M., Zhang, L., Kumar, A. A., & Jones, M. N. (2022). forager: A Python framework for Modeling Mental Search. *Paper presented at the Annual Meeting of the Society for Computation in Psychology*.

Aslan, A., Bäuml, K.-H., & Grundgeiger, T. (2007). The role of inhibitory processes in part-list cuing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*(2), 335.

Auriacombe, S., Lechevallier, N., Amieva, H., Harston, S., Raoux, N., & Dartigues, J.-F. (2006). A longitudinal study of quantitative and qualitative features of category verbal fluency in incident Alzheimer's disease subjects: results from the PAQUID study. *Dementia and geriatric cognitive disorders, 21*(4), 260-266.

Barber, S. J., Harris, C. B., & Rajaram, S. (2015). Why two heads apart are better than two heads together: multiple mechanisms underlie the collaborative inhibition effect in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 41*(2), 559-566.

Barber, S. J., & Rajaram, S. (2011). Exploring the relationship between retrieval disruption from collaboration and recall. *Memory, 19*(5), 462-469.

Barber, S. J., Rajaram, S., & Aron, A. (2010). When two is too many: Collaborative encoding impairs memory. *Memory & Cognition, 38*(3), 255-264.

Barber, S. J., Rajaram, S., & Fox, E. B. (2012). Learning and remembering with others: The key role of retrieval in shaping group recall and collective memory. *Social Cognition, 30*(1), 121.

Basden, B. H., Basden, D. R., Bryner, S., & Thomas, R. L. (1997). A comparison of group and individual remembering: Does collaboration disrupt retrieval strategies? *Journal of Experimental Psychology-Learning Memory and Cognition, 23*(5), 1176-1191. doi:Doi 10.1037/0278-7393.23.5.1176

Basden, B. H., Basden, D. R., & Henry, S. (2000). Costs and benefits of collaborative remembering. *Applied Cognitive Psychology, 14*(6), 497-507. doi:Doi 10.1002/1099-0720(200011/12)14:6<497::Aid-Acp665>3.0.Co;2-4

Basden, D. R., & Basden, B. H. (1995). Some tests of the strategy disruption interpretation of part-list cuing inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 1656-1669.

Basden, D. R., & Draper, J. S. (1973). Effect of cuing with list members as a function of number of categories, taxonomic frequency, and presentation order. *Canadian Journal of Psychology*, *27*, 327-333.

Bäuml, K.-H., & Aslan, A. (2004). Part-list cuing as instructed retrieval inhibition. *Memory & Cognition, 32*(4), 610-617.

Bäuml, K.-H., & Aslan, A. (2006). Part-list cuing can be transient and lasting: the role of encoding. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*(1), 33.

Blumen, H. M., & Rajaram, S. (2008). Influence of re-exposure and retrieval disruption during group collaboration on later individual recall. *Memory, 16*, 231-244.

Bouchard, T. J., Jr., & Hare, M. (1970). Size, performance, and potential in brainstorming groups. *Journal of Applied Psychology, 54*, 51-55.

Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical population biology, 9*(2), 129-136.

Choi, H.-Y., Blumen, H. M., Congleton, A. R., & Rajaram, S. (2014). The role of group configuration in the social transmission of memory: Evidence from identical and reconfigured groups. *Journal of Cognitive Psychology, 26*(1), 65-80.

Choi, H.-Y., Kensinger, E. A., & Rajaram, S. (2017). Mnemonic transmission, social contagion, and emergence of collective memory: Influence of emotional valence, group structure, and information distribution. *Journal of Experimental Psychology: General, 146*(9), 1247.

Clark, S. E., Hori, A., Putnam, A., & Martin, T. P. (2000). Group collaboration in recognition memory. *Journal of Experimental Psychology-Learning Memory and Cognition, 26*(6), 1578-1588. doi:Doi 10.1037/0278-7393.26.6.1578

Congleton, A. R., & Rajaram, S. (2011). The influence of learning methods on collaboration: prior repeated retrieval enhances retrieval organization, abolishes collaborative inhibition, and promotes post-collaborative memory. *Journal of Experimental Psychology: General, 140*(4), 535.

Congleton, A. R., & Rajaram, S. (2014). Collaboration changes both the content and the structure of memory: Building the architecture of shared representations. *Journal of Experimental Psychology: General, 143*(4), 1570.

Dahlström, Ö., Danielsson, H., Emilsson, M., & Andersson, J. (2011). Does retrieval strategy disruption cause general and specific collaborative inhibition? *Memory, 19*(2), 140-154.

Danielsson, H., Dahlström, Ö., & Andersson, J. (2011). The more you remember the more you decide: Collaborative memory in adolescents with intellectual disability and their assistants. *Research in Developmental Disabilities, 32*(2), 470-476.

Diehl, M., & Stroebe, W. (1987). Productivity Loss in Brainstorming Groups - toward the Solution of a Riddle. *Journal of Personality and Social Psychology, 53*(3), 497-509. doi:Doi 10.1037/0022-3514.53.3.497

Finlay, F., Hitch, G. J., & Meudell, P. R. (2000). Mutual inhibition in collaborative recall: Evidence for a retrieval-based account. *Journal of Experimental Psychology-Learning Memory and Cognition, 26*(6), 1556-1567. doi:10.1037//0278-7393.26.6.1556

Gates, V., Suchow, J. W., & Griffiths, T. L. (2022). Memory transmission in small groups and large networks: An empirical study. *Psychonomic Bulletin & Review, 29*(2), 581-588.

Harris, C. B., Barnier, A. J., & Sutton, J. (2012). Consensus collaboration enhances group and individual recall accuracy. *Quarterly Journal of Experimental Psychology, 65*(1), 179-194.

Harris, C. B., Barnier, A. J., Sutton, J., & Savage, G. (2019). Features of successful and unsuccessful collaborative memory conversations in long-married couples. *Topics in Cognitive Science, 11*(4), 668-686.

Harris, C. B., Paterson, H. M., & Kemp, R. I. (2008). Collaborative recall and collective memory: what happens when we remember together? *Memory, 16*(3), 213-230. doi:10.1080/09658210701811862

He, L., Richie, R., & Bhatia, S. (2022). *Limitations to Optimal Search in Naturalistic Active Learning.* Paper presented at the Proceedings of the Annual Meeting of the Cognitive Science Society.

Henkel, L. A., & Rajaram, S. (2011). Collaborative remembering in older adults: Age-invariant outcomes in the context of episodic recall deficits. . *Psychology and Aging, 26*(3), 532-545.

Henry, J. D., & Crawford, J. R. (2004). Verbal fluency deficits in Parkinson's disease: a meta-analysis. *Journal of the International Neuropsychological Society, 10*(4), 608-622.

Herzmann, G., & Curran, T. (2011). Experts' memory: an ERP study of perceptual expertise effects on encoding and recognition. *Memory & Cognition, 39*(3), 412-432.

Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science, 30*(1), 3-41.

Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review, 119*(2), 431.

Hirst, W., & Manier, D. (2008). Towards a psychology of collective memory. *Memory, 16*(3), 183-200.

Hoppe, R. A. (1962). Memorizing by individuals and groups. *Journal of Abnormal and Social Psychology, 65*, 64-67.

Howard, M. W., Jing, B., Addis, K. M., & Kahana, M. J. (2007). Semantic structure and episodic memory. In *Handbook of latent semantic analysis* (pp. 133-154): Psychology Press.

Hultsch, D. F. (1971). Adult age differences in free classification and free recall. *Developmental Psychology, 4*(3), 338-342.

Hutchinson, J. M., Wilke, A., & Todd, P. M. (2008). Patch leaving in humans: can a generalist adapt its rules to dispersal of items across patches? *Animal Behaviour, 75*(4), 1331-1349.

Ingham, A. G., Levinger, G., Graves, J., & Peckham, V. (1974). The Ringelmann effect: Studies of group size and group performance. *Journal of Experimental Social Psychology, 10*(4), 371-384.

Johansson, O., Andersson, J., & Ronnberg, J. (2000). Do elderly couples have a better prospective memory than other elderly people when they collaborate? *Applied Cognitive Psychology*, *14*, 121-133.

Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, *114*(1), 1-37.

Kelley, M. R., Reysen, M. B., Ahlstrand, K. M., & Pentz, C. J. (2012). Collaborative inhibition persists following social processing. *Journal of Cognitive Psychology, 24*(6), 727-734.

Latane, B., & Nida, S. (1981). Ten years of research on group size and helping. *Psychological Bulletin, 89*(2), 308.

Latane, B., Williams, K., & Harkins, S. (1979). Many hands make light the work: The causes and consequences of social loafing. *Journal of Personality and Social Psychology, 37*, 822-832.

Lehmer, E.-M., & Bäuml, K.-H. T. (2018). Part-list cuing can impair, improve, or not influence recall performance: The critical roles of encoding and access to study context at test. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(8), 1186.

Lorge, J., & Solomon, H. (1955). Two models of group behavior in the solution of eureka-type problems. *Psychometrika, 20*, 139-148.

Luhmann, C. C., & Rajaram, S. (2015). Memory transmission in small groups and large networks: An agent-based model. *Psychological Science*, *26*(12), 1909-1917.

Lundin, N. B. (2022). Disorganized speech in psychosis: Computational and neural markers of semantic foraging and discourse cohesion. *Doctoral Dissertation, Indiana University, ProQuest Dissertations Theses Global*.

Lundin, N. B., Todd, P. M., Jones, M. N., Avery, J. E., O'Donnell, B. F., & Hetrick, W. P. (2020). Semantic search in psychosis: Modeling local exploitation and global exploration. *Schizophrenia bulletin open, 1*(1), sgaa011.

Mantyla, T., & Nilsson, L.-G. (1983). Are my cues better than your cues?| Uniqueness and reconstruction as prerequisites for optimal recall of verbal materials. *Scandinavian Journal of Psychology, 24*(1), 303-312.

Maswood, R., & Rajaram, S. (2019). Social transmission of false memory in small groups and large networks. *Topics in Cognitive Science, 11*(4), 687-709.

Meade, M. L., & Gigone, D. (2011). The effect of information distribution on collaborative inhibition. *Memory, 19*(5), 417-428.

Meade, M. L., Nokes, T. J., & Morrow, D. G. (2009). Expertise promotes facilitation on a collaborative memory task. *Memory, 17*(1), 39-48. doi:Pii 907094176

10.1080/09658210802524240

Meade, M. L., & Roediger, H. L. (2009). Age differences in collaborative memory: The role of retrieval manipulations. *Memory & Cognition, 37*(7), 962-975.

Meudell, P. R., Hitch, G. J., & Boyle, M. M. (1995). Collaboration in recall: Do pairs of people cross cue each other to produce new memories? *Quarterly Journal of Experimental Psychology, 48A*, 141-152.

Meudell, P. R., Hitch, G. J., & Kirby, P. (1992). Are two heads better than one? Experimental investigations of the social facilitaion of memory. *Applied Cognitive Psychology, 6*, 525-543.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in information processing systems*, 3111-3119.

Miller, B., & Steyvers, M. (2011). *The wisdom of crowds with communication.* Paper presented at the Proceedings of the annual meeting of the cognitive science society.

Monsch, A. U., Bondi, M. W., Butters, N., Salmon, D. P., Katzman, R., & Thal, L. J. (1992). Comparisons of verbal fluency tasks in the detection of dementia of the Alzheimer type. *Archives of neurology, 49*(12), 1253-1258.

Morrow, D. G., Miller, L. M., Ridolfo, H. E., Magnor, C., Fischer, U. M., Kokayeff, N. K., & Stine-Morrow, E. A. (2008). Expertise and age differences in pilot decision making. *Aging, Neuropsychology, and Cognition, 16*(1), 33-55.

Nickerson, R. S. (1984). Retrieval Inhibition from Part-Set Cueing - a Persisting Enigma in Memory Research. *Memory & Cognition, 12*(6), 531-552. doi:Doi 10.3758/Bf03213342

Paulus, P. B., & Dzindolet, M. T. (1993). Social influence processes in group brainstorming. *Journal of Personality and Social Psychology, 64*, 575-586.

Perlmutter, H. V. (1953). Group memory of meaningful material. *Journal of Abnormal and Social Psychology, 47*(762-769).

Perlmutter, M. (1979). Age differences in adults' free recall, cued recall, and recognition. *Journal of Gerontology, 34*(4), 533-539.

Piaget, J. (1959). *The language and thought of the child.* Thetford, Norfolk: Lowe & Brydone Printers Ltd.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. *The psychology of learning and motivation: Advances in research and theory, 14*, 207-262.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88,* 93-134.

Rajaram, S., & Pereira-Pasarin, L. P. (2010). Collaborative Memory: Cognitive Research and Theory. *Perspect Psychol Sci, 5*(6), 649-663. doi:10.1177/1745691610388763

Roediger, H. L. (1978). Recall as a self-limiting process. *Memory & Cognition, 6*, 54-63.

Rundus, D. (1973). Negative effects of using list items as recall cues. *Journal of Verbal Learning and Verbal Behavior*, *12*(1), 43-50.

Ryack, B. L. (1965). A comparison of individual and group learning of nonsense syllables. *Journal of Personality and Social Psychology, 2*, 296-299.

Slamecka, N. J. (1968). An Examination of Trace Storage in Free Recall. *Journal of Experimental Psychology, 76*(4p1), 504-&. doi:DOI 10.1037/h0025695

Stephenson, G. M., Abrams, D., Wagner, W., & Wade, G. (1986). Partners in recall: Collaborative order in the recall of police interrogation. *British Journal of Social Psychology, 25*, 341-343.

Steyvers, M., Miller, B., Hemmer, P., & Lee, M. (2009). The wisdom of crowds in the recollection of order information. *Advances in neural information processing systems, 22*.

Surowiecki, J. (2005). *The wisdom of crowds*: Anchor.

Taylor, D. W., Berry, P. C., & Block, C. H. (1958). Does group participation when using brainstorming facilitate or inhibit creative thinking? *Administrative Science Quarterly, 3*, 23-47.

Thorley, C., & Dewhurst, S. A. (2007). Collaborative false recall in the DRM procedure: Effects of group size and group pressure. *European Journal of Cognitive Psychology*, *19*(6), 867-881.

Thorley, C., & Dewhurst, S. A. (2009). False and veridical collaborative recognition. *Memory, 17*(1), 17-25.

Troyer, A. K., Moscovitch, M., & Winocur, G. (1997). Clustering and switching as two components of verbal fluency: evidence from younger and older healthy adults. *neuropsychology, 11*(1), 138.

Troyer, A. K., Moscovitch, M., Winocur, G., Leach, L., & Freedman, M. (1998). Clustering and switching on verbal fluency tests in Alzheimer's and Parkinson's disease. *Journal of the International Neuropsychological Society*, *4*(2), 137-143.

Tulving, E. (1974). Cue-dependent forgetting. *American Scientist, 62*, 74-82.

Wang, Q. (2008). On the cultural constitution of collective memory. *Memory, 16*, 305-317.

Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In *Theories of group behavior* (pp. 185-208). New York, NY: Springer.

Wegner, D. M., Erber, R., & Raymond, P. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology, 61*(6), 923.

Weldon, M. S. (2000). Remembering as a social process. In *Psychology of learning and motivation* (Vol. 40, pp. 67-120): Elsevier.

Weldon, M. S., & Bellinger, K. D. (1997). Collective memory: Collaborative and individual processes in remembering. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 1160–1175.

Weldon, M. S., Blair, C., & Huebsch, D. (2000). Group remembering: Does social loafing underlie collaborative inhibition? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 1568-1577.

Wilke, A., Hutchinson, J. M., Todd, P. M., & Czienskowski, U. (2009). Fishing for the right words: Decision rules for human foraging behavior in internal search tasks. *Cognitive Science, 33*(3), 497-529.

Wilson, J. H., & Criss, A. H. (2017). The list strength effect in cued recall. *Journal of Memory and Language*, *95*, 78-88.

Wilson, J. H., Kellen, D., & Criss, A. H. (2020). Mechanisms of output interference in cued recall. *Memory & Cognition, 48*(1), 51-68.

Wolfe, J. M. (2013). When is it time to move to the next raspberry bush? Foraging rules in human visual search. *Journal of vision, 13*(3), 10-10.

Wright, D. B., & Klumpp, A. (2004). Collaborative inhibition is due to the product, not the process, of recalling in groups. *Psychonomic Bulletin & Review, 11*(6), 1080-1083.

Yaron-Antar, A., & Nachson, I. (2006). Collaborative remembering of emotional events: the case of Rabin's assassination. *Memory, 14*(1), 46-56. doi:10.1080/09658210444000502

Willa M. Mannering

EDUCATION

PhD Double major in Cognitive Science and Psychology **June 2023**
Indiana University, Bloomington, Indiana

BS Cognitive Science **May 2018**
Indiana University, Bloomington, Indiana
**With Highest Distinction*

BA Computer Science **May 2018**
Indiana University, Bloomington, Indiana
**With Highest Distinction*

RESEARCH EXPERIENCE

2022—Present

Neurosymbolic AI Graduate Intern
National Renewable Energy Lab, Golden, Co
Topics: renewable energy system management, neurosymbolic AI algorithms, planning algorithms

- Developing human interpretable neurosymbolic AI algorithms for planning and managing energy systems.

2018—2023

Doctoral Researcher
Cognitive Computing Laboratory, Indiana University, Department of Psychological and Brain Sciences
Topics: catastrophic interference, semantic models of memory, semantic search, collaborative memory

- Conducted experiment and analyzed behavioral data to determine the effect of group dynamics on idea production.
- Designed and implemented cognitive architectures in Python to model and make actionable predictions about how people remember and learn in groups.
  - Commonly used popular Python libraries such as NumPy, SciPy, and Pandas.
- Designed and conducted a 2-year research project to study the effect of catastrophic forgetting in large language models.

2014—2018

Undergraduate Research Assistant
Visual Perception and Electrophysiological Laboratory, Indiana University, Department of Psychological and Brain Sciences
Topics: latent fingerprint examination, taboo tradeoffs

- Designed and implemented a multi-year research project with the goal of comparing the decision criteria for fingerprint examiners and novices during fingerprint examinations.

- Developed online web visualization and data collection protocol via Amazon Mechanical Turk
- Analyzed results using advanced regression models

2016—2018 **Undergraduate Research Assistant**
Percepts and Concepts Laboratory, Indiana University, Department of Psychological and Brain Sciences
Topics: face categorization

- Conducted a 1-year research project on face perception and categorization
  - Created all stimuli, collected and analyzed resulting experimental data

TEACHING EXPERIENCE

2018—2023 **Associate Instructor**
Indiana University, Statistical Techniques
Indiana University, Cognitive Psychology
Indiana University, Methods of Experimental Psychology
Indiana University, Human Memory

2016 **Undergraduate Teaching Intern**
Indiana University, Foundations of Cognitive Science Course

PUBLICATIONS

1. **Mannering, W. M.**, Todd, P. M. & Jones, M. N. (In Progress) Optimal Foraging Behaviors in Collaborative Verbal Fluency Tasks.
2. **Mannering, W. M.**, Rajaram, S., Shiffrin, R. M. & Jones, M. N. (2023) Modeling Collaborative Memory with SAM. *Memory and Cognition.* Under Review.
3. Krendl, A. C., **Mannering, W. M.,** Jones, M. N., Hugenberg, K., & Kennedy, K. P. (2022) Determining whether older adults use similar strategies to young adults in theory of mind tasks. *Journal of Gerontology: Psychological Sciences.*
4. **Mannering, W. M.,** Rajaram, S., Shiffrin, R. M., & Jones, M. N. (2022) Modeling the Effect of Learning During Retrieval on Collaborative Inhibition. Proceedings of the 44th Annual Meeting of the Cognitive Science Society.
5. **Mannering, W. M.**, Rajaram, S., & Jones, M. N. (2021) Towards a Cognitive Model of Collaborative Memory. Proceedings of the 43rd Annual Meeting of the Cognitive Science Society. 959-965.

6. **Mannering, W. M.**, Vogelsang, M., Busey, T. & Mannering, F. (2021) Are Forensic Scientists Too Risk Averse? *Journal of Forensic Sciences*. 66(4), 1377-1400. https://doi.org/10.1111/1556-4029.14700

7. **Mannering, W. M.** & Jones, M. N. (2020) Catastrophic Interference in Neural Network Models of Distributional Semantics. *Computational Brain & Behavior*. https://doi.org/10.1007/s42113-020-00089-5

8. **Mannering, W. M.** & Jones, M. N. (2019) Insulating Distributional Semantic Models from Catastrophic Interference. Proceedings of the 41st Annual Meeting of the Cognitive Science Society. 2282-2288.

9. **Mannering, W. M.,** Vogelsang, M., & Busey, T. (2017). Should Fingerprint Examiners Make More Erroneous Identifications? *Indiana University of Undergraduate Research*, *3*(1), 68-77.

PRESENTATIONS

1. **Mannering, W. M.**, Rajaram, S., Shiffrin, R. M., & Jones, M. N. (2022) Towards a Cognitive Model of Collaborative Memory. Accepted as poster at the 44th Annual Meeting of the Cognitive Science Society.
2. **Mannering, W. M.**, Rajaram, S., & Jones, M. N. (2021) Towards a Cognitive Model of Collaborative Memory. Presented at the Ninth Annual Conference on Advances in Cognitive Systems

3. **Mannering, W. M.**, Rajaram, S., & Jones, M. N. (2021) Towards a Cognitive Model of Collaborative Memory. Presented at the 43rd Annual Meeting of the Cognitive Science Society.

4. **Mannering, W. M.** & Jones, M. N. (2019) Insulating Distributional Semantic Models from Catastrophic Interference. Presented as poster at the 60th Annual Meeting of the Psychonomic Society.

5. **Mannering, W. M.** & Jones, M. N. (2019) Insulating Distributional Semantic Models from Catastrophic Interference. Presented at the 8th Annual Midwest Cognitive Science Conference.

6. **Mannering, W. M.** & Jones, M. N. (2019) Insulating Distributional Semantic Models from Catastrophic Interference. Presented as poster at the 41st Annual Meeting of the Cognitive Science Society.

7. **Mannering, W. M.,** Vogelsang, M., Busey, T. (2017). Should Fingerprint Examiners Make More Erroneous Identifications? Presented as poster at the 56th Annual Meeting of the Psychonomic Society.

8. **Mannering, W. M.,** Vogelsang, M., Busey, T. (2016). Should Fingerprint Examiners Make More Erroneous Identifications? Presented as poster at the 8th Annual Midwest Undergraduate Cognitive Science Conference.

ACADEMIC SERVICE

2023 **Committee Member**

Standing Committee on Statistical Methods for the Transportation Research Board of the National Academies of Sciences, Engineering, and Medicine

2016—2018 **Director**
Midwest Undergraduate Cognitive Science Conference (MUCSC)

AWARDS AND ACHIEVEMENTS

2019 IUB Provost's Travel Award for Women in Science

2019 2019 Sharon Stephens Brehm Excellence in Research Award from the IU Department of Psychological and Brain Sciences

2019 Outstanding Teaching Assistant Award from the IU Department of Psychological and Brain Sciences

2019 Honorable Mention for NSF Graduate Research Fellowship Program

2018 Phi Beta Kappa Honor Society Membership

2018 Outstanding Undergraduate Research Award from the IU Cognitive Science Program

VOLUNTEER WORK

2022—2023 Hack for LA Data Scientist Volunteer

2021 COVID Vaccine Clinic volunteer, volunteered with Monroe County Health Department to help administer COVID-19 vaccines

2017—2018 WonderLab Museum of Science, general volunteer work monitoring young guests in the museum

2015 Banneker Buddies, volunteered as a tutor for in need elementary students at Fairview Elementary School