Modeling the Effect of Learning During Retrieval on Collaborative Inhibition

Willa M. Mannering[1], Suparna Rajaram[2], Richard M. Shiffrin[1], Michael N. Jones[1]
[1]Indiana University, Bloomington
[wmanneri][shiffrin][jonesmn]@iu.edu

[2]Stony Brook University, Stony Brook
[suparna.rajaram]@stonybrook.edu

Abstract

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 comparable size. Previously, we adapted the Search of Associative Memory (SAM; Raaijmakers & Shiffrin, 1981) model to collaborative free recall and found support for the retrieval disruption hypothesis as the cause of collaborative inhibition (Mannering, Rajaram, Shiffrin, & Jones, 2022). In this paper, we investigate another possible cause of collaborative inhibition: memory homogenization. Previous modeling attempts of collaborative recall have shown that memory homogenization may contribute to collaborative inhibition (Luhmann & Rajaram, 2015). To determine the effect of memory homogenization and retrieval disruption on collaborative inhibition in SAM, we prevent SAM models from learning during collaborative retrieval which subsequently prevents memory homogenization. We found that even when SAM model memories remained diversified, collaborative inhibition persisted—though the strength of the effect was diminished. These results suggest that both retrieval disruption and memory homogenization may contribute to collaborative inhibition in the SAM model.

Keywords: collaborative memory, collaborative inhibition, memory modeling, search of associative memory

Introduction

Humans usually form and retrieve memories in groups. However, most of our knowledge of human memory comes from studying individuals alone. The field of collaborative memory aims to remedy this discrepancy by taking a cognitive approach to studying group memory. *Collaborative inhibition* is the primary phenomenon studied within this field and is the tendency for collaborative groups to underperform in free recall tasks compared to nominal groups of comparable size (Basden, Basden, Bryner, & Thomas, 1997; Weldon & Bellinger, 1997). Collaborative inhibition has been widely studied within the field of collaborative memory and is found in many diverse group compositions and materials (Andersson & Ronnberg, 1995; Marion & Thorley, 2016; Meudell, Hitch, & Kirby, 1992; Rajaram & Pereira-Pasarin, 2010).

Within the literature, there are several competing hypotheses for collaborative inhibition. The explanation with the most empirical support is the retrieval disruption hypothesis which posits that collaborative inhibition occurs because individual retrieval strategies are disrupted during group recall (Basden et al., 1997). As of now, most of the research within the collaborative memory field is behavioral. However, there are limitations to behavioral research that can and should be supplemented by modeling efforts.

Modeling Collaborative Inhibition

Behavioral researchers are unable to comprehensively study the internal mechanisms at play during collaborative retrieval. While there is an abundance of behavioral support for the retrieval disruption hypothesis, until recently, there were no cognitive models that could provide a closer look at what might be happening cognitively during collaborative retrieval.

Search of Associative Memory (SAM) Recently, we adapted the well-validated Search of Associative Memory (SAM; Raaijmakers & Shiffrin, 1981) model to collaborative free recall. SAM is a cue-dependent probabilistic search theory of retrieval of free recall and free recall with cues. It has two phases, encoding and retrieval, and two memory systems, short-term memory and long-term memory. The model begins by encoding a list of study items. In this phase, it uses a buffer rehearsal system so that items that appear together in the short-term buffer have higher associations with each other when transferred to long-term memory. Long-term memory is where information is transferred from short-term memory and stored permanently. Long-term memory storage has two components: an association vector of study items to environmental context and an association matrix of study item to study item information. Retrieval in SAM begins by using cues from short-term memory as probes for long-term memory. These cues include context cues (context recall) and previously recalled words from the study list (word cue recall). Equation 1a gives the probability of sampling a word, W_{iS}, using only context, C_T, as a memory probe. Equation 1b gives the probability of sampling a word, W_{iS}, given both context, C_T, and a word cue, W_{kT}, as a memory probe. The T subscript is used to indicate a cue at test and the S subscript indicates the item as it is stored in memory.

Table 1. SAM Parameter Descriptions

Parameter Name	Description
t	Presentation time per word during encoding
r	STM buffer size
a	Weight for context to word association during encoding
b	Weight for word to other word association during encoding
c	Weight for word cue to same word association during encoding
d	Residual strength of association for words that never appear in buffer together during encoding
e	Incrementing parameter for context-to-word association during retrieval
f	Incrementing parameter for word-to-word association during retrieval
g	Incrementing parameter for word-to-self association during retrieval
Kmax	Maximum number of retrieval failures before retrieval is stopped
Lmax	Maximum number of retrieval attempts using word cues before returning to context cues

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

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

Table 1 gives a brief description of the standard parameters included in the model. The original SAM model accounts for a broad range of free recall phenomena including serial position curves, presentation duration effects, list length effects, extended recall, and repeated recall.

Collaborative SAM To adapt SAM to collaborative free recall we created a shared memory buffer between two or more models which represents words "spoken" aloud by the models. The models begin retrieval by performing context recall separately. The first response produced by any of the models in the group is added to the shared buffer. Once a response is in the shared buffer, the other models in the group can access the response. An extra parameter, *j,* was added as an incrementing parameter between words added to the group response. This represents learning of the retrieval organization created by the group during recall. Then, all the models use the new response in the shared buffer as a cue to perform recall. The models perform cued recall separately and the response produced first is added to the shared buffer. The models continue using new responses as cues until no models can produce any new responses. At this point the models return to context recall. The process of switching between context and cued recall continues until all the models reach a specified number of retrieval failures at which point the retrieval phase ends.

We found that the collaborative SAM model produced the basic patterns of collaborative inhibition seen in the experimental literature in both categorized and uncategorized word lists. We also found some preliminary evidence that the collaborative SAM model supports the retrieval disruption hypothesis while fitting the model to categorized list data from the Basden et al. (1997) study. Basden et al. (1997) found that collaborative inhibition is stronger when category size is larger and study materials are less organized. When category sizes are larger, there is more room for more idiosyncratic organization and retrieval strategies within individual categories between group members. Conversely, when category sizes are smaller, there is less room for idiosyncratic organization and retrieval strategies within individual categories between group members. When the internal organization of study items is dissimilar between group members (which is more likely in the larger category condition), collaborative inhibition increases because the cues from other group members tend to disrupt individual search strategies. We found that SAM reproduced this effect as collaborative inhibition increased when the category size was large compared to when the category size was small. Reproducing this effect suggests that SAM supports the retrieval disruption hypothesis.

To investigate how SAM might be producing the collaborative inhibition effect, we fit the model to experimental data (from Choi, Blumen, Congleton, & Rajaram, 2014) and estimated 5 relevant parameters. We found that the two stopping parameters, *Kmax* and *Lmax* were significantly different between the collaborative and nominal groups, with collaborative groups having a lower *Kmax* value and a higher *Lmax* value than nominal groups. These findings suggest that SAM may produce collaborative inhibition by having collaborative groups recall for less time overall than the nominal groups while also spending more time on each word cue than nominal groups. The goal of this paper is to further evaluate the retrieval disruption hypothesis by considering the impact of other mechanisms that could potentially produce collaborative inhibition in SAM.

Previous Modeling Attempts Before our collaborative SAM model, 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 memory, collaborative inhibition was seen in the recall of groups of 3 agents. Additionally, they were able to model some predictions of the collaborative memory field, namely the effect of group size on collaborative inhibition. However, while this model included psychologically based agents that were able to encode and retrieve memories, the implementation was not as mechanistically extensive as SAM 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. The learning that the agents achieve during the collaborative recall task decreases the diversity of the 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.

In this paper, our goal is to determine whether memory homogenization could be a cause of collaborative inhibition in SAM. Previously, our collaborative SAM model showed support for the retrieval disruption hypothesis, but the results of the Luhmann and Rajaram (2015) study suggest there may be other causes to consider. To tease apart the underlying cause of collaborative inhibition in SAM, we plan to prevent memory homogenization during collaborative recall by preventing learning during the retrieval phase in the models. If collaborative inhibition is still present when the model memories remain diverse, then we can rule out memory homogenization as the sole cause of collaborative inhibition in SAM.

Part 1: Learning in SAM

Learning during recall is a natural characteristic of the individual version of SAM. The 3 parameters responsible for learning in SAM are the *e* (the incrementing parameter for context-to-word association), *f* (the incrementing parameter for word-to-word association), and *g* (the incrementing parameter for word-to-self association) parameters.

Memory Structures in SAM

The model has two stages of memory: short-term and long-term. Short-term memory is where encoding and rehearsal occur, achieved with a buffer system, while long term memory consists of context memory and word association memory. The context memory is a vector containing context and study item associations. The strength of the context associations is dependent on the time a study item spends in the short-term memory buffer.

The word association memory is a matrix containing study item to study item associations. The strength of the association between words is dependent on the amount of time two words spent together in the short-term memory buffer. For words that never appear together in the buffer, there is a residual association strength assigned based on the value of parameter *d.* Figure 1 shows the two forms of long-term memory in the SAM model.

Context Association Vector

$$\begin{array}{cc} & \begin{array}{ccc} w_1 & w_2 & w_3 \end{array} \\ \text{context} & \begin{pmatrix} w_{1c} & w_{2c} & w_{3c} \end{pmatrix} \end{array}$$

Word Association Matrix

$$\begin{array}{cc} & \begin{array}{ccc} w_1 & w_2 & w_3 \end{array} \\ \begin{array}{c} w_1 \\ w_2 \\ w_3 \end{array} & \begin{pmatrix} (w_1, w_1) & (w_1, w_2) & (w_1, w_3) \\ (w_2, w_1) & (w_2, w_2) & (w_2, w_3) \\ (w_3, w_1) & (w_3, w_2) & (w_3, w_3) \end{pmatrix} \end{array}$$

Figure 1. Long term memory structures in SAM. The context association vector represents the associations between study items and context. The word association matrix represents associations between study items.

Learning During Individual Retrieval

The first phase of retrieval in SAM uses context memory as a starting point. Words that are likely to be recalled in this phase have a high context association. When a word, W_1, is recalled in this phase of retrieval, the context memory is updated according to *e* and the W_1 association to itself is updated according to *g.* The second phase of retrieval uses previously recalled words as cues to recall more words. In this case, W_1 is used as a cue and words that are likely to be recalled have a high word association with W_1. When a new word is recalled in this phase of retrieval, the context memory is updated according to *e*, the W_1 association to itself is updated according to *g,* and the association between W_1 and W_2 is updated according to *f.* Thus, as SAM recalls words during retrieval, the associations between cues and recalled words becomes stronger—producing the effect of learning over time. Figure 2 is a visualization of how learning occurs in SAM during the retrieval phase.

Learning Allowed During Collaborative Retrieval

Both the individual and collaborative versions of SAM are able to learn during retrieval. Learning in a collaborative group of SAM models is the same as individual models with the addition of one parameter, *j.* Parameter *j* is the incrementing parameter for a cue word to the group response and controls how the group response influences models' memories during retrieval.

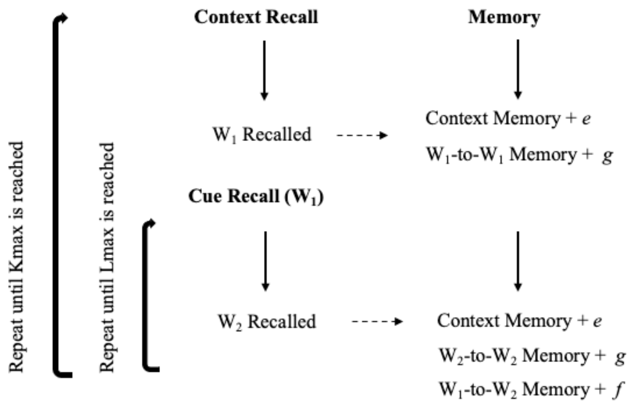

Figure 2. Diagram of Learning during Retrieval in SAM.

While learning does occur during collaborative retrieval, it has not yet been shown that model memories homogenize when retrieving in groups. To measure the similarity of model memories during retrieval, we recorded the overall cosine similarity between each model's word association memories. The cosine similarity between two vectors is their dot product divided by the product of their magnitudes.

To measure the change in cosine similarity between multiple models' memories we compared the individual word representations within each model at 20 different timestamps during retrieval. The timestamps used in these calculations are the points during retrieval when a word is successfully recalled by the group and the models' memories are updated (see Figure 2). This process works as follows: if there are 2 study words on the study list and 2 models in a group, then we would find the mean cosine similarity of word 1 (represented by Row 1 in Figure 3) between models 1 and 2 and word 2 (represented by Row 2 in Figure 3) between models 1 and 2 for timestamp 1. Then we would continue this process for all 20 timestamps. For a visualization of this process over one timestamp, see Figure 3.

The result of this process is Figure 4 which shows the average cosine similarity (over 200 collaborative retrieval simulations) of the word association memories of a model group over 20 timestamps during retrieval. Over the course of retrieval, the association memories of models in a collaborative group do become more similar to each other (homogenization). This is consistent with the finding that model memories homogenize and become less diverse over the course of collaborative retrieval (Luhmann & Rajaram, 2015).

Part 2: Cause of Collaborative Inhibition During Retrieval

In the previous section, we found that the collaborative SAM models' memories homogenized over the course of retrieval. Additionally, we have previously shown that these models produce collaborative inhibition (Mannering et al., 2022). The next step in this study is to determine whether learning during retrieval, which causes model memories to homogenize and become less diverse, is contributing to collaborative inhibition in the SAM models.

Learning Prevented During Collaborative Retrieval

To determine whether homogenized memories cause collaborative inhibition, we prevented learning during retrieval—something that is clearly not possible in behavioral experiments. If no learning occurs during retrieval, then the model memories will not homogenize and become less diverse. If collaborative inhibition persists, then this supports the claim that model memory homogenization is not the sole cause of collaborative inhibition in the collaborative SAM models.

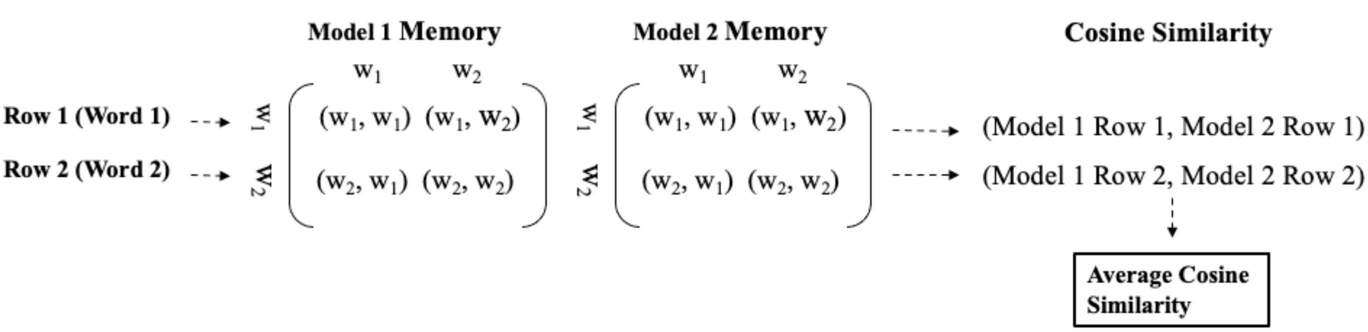

Figure 3. Cosine similarity calculation during one timestamp in collaborative retrieval. Row 1 represents the association vectors for word 1 in Model 1 and Model 2, respectively. The cosine similarity of these two vectors is calculated and recorded. Row 2 represents the association vectors for word 2 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.

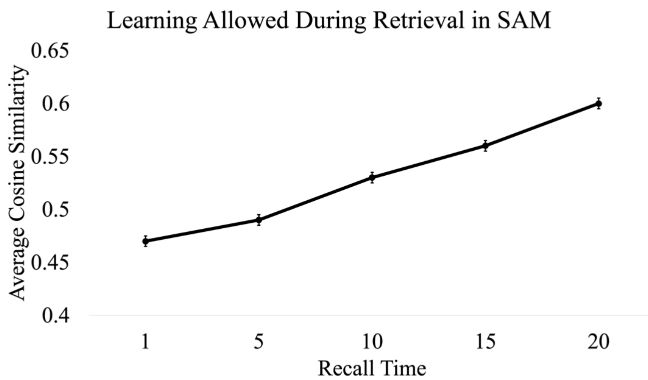

Figure 4. The mean cosine similarity (over 200 collaborative retrieval simulations) of models' associative memories during collaborative retrieval over 20 timestamps.

To accomplish this task, we set the parameters responsible for learning during recall in the collaborative models (*e, f, g,* and *j*) to 0 so that the model does 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, we then repeated the method of measuring model memory similarity over retrieval detailed in the previous section of this paper (see Equation 1 and Figure 3). Figure 5 shows the change in models' memory similarity over retrieval when learning was prevented.

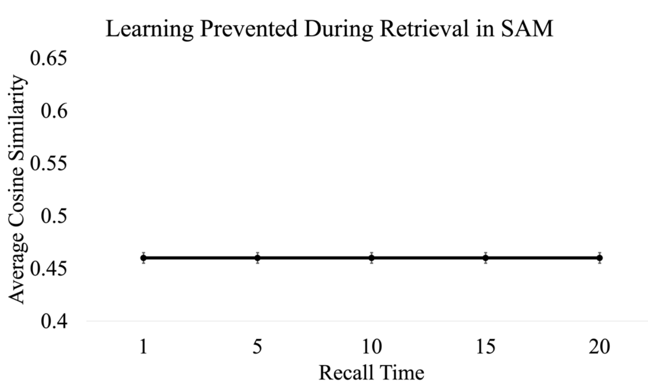

Figure 5. The mean cosine similarity (over 200 collaborative retrieval simulations) of models' associative memories during collaborative retrieval over 20 timestamps when learning was prevented.

As depicted in Figure 5, when learning is prevented, models' 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.

After we determined that preventing learning during retrieval also prevents model memory homogenization, we evaluated how collaborative inhibition was affected during collaborative retrieval. In Figure 6, we compared the amount of collaborative inhibition in the learning allowed and learning prevented conditions of collaborative retrieval. Figure 6 shows that collaborative inhibition persists in the learning prevented condition where the models' memories do not homogenize over retrieval, however, the size of the inhibitory effect is diminished. This suggests that collaborative inhibition in collaborative SAM models is not caused solely by increased memory homogeneity over retrieval. Additionally, in Figure 6, the overall proportion recalled was higher in the learning prevented condition than in the learning allowed condition.

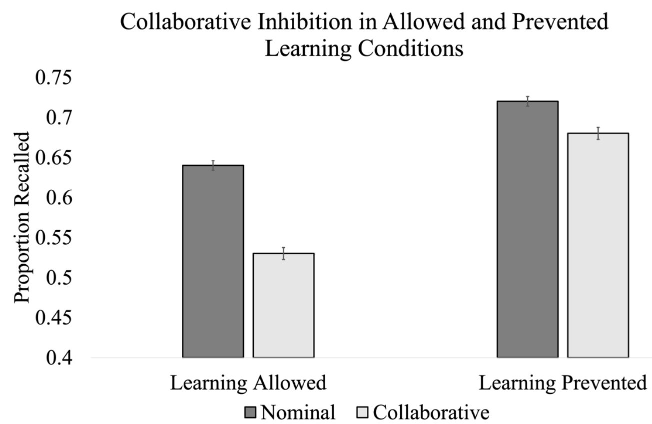

Figure 6. Comparison of collaborative inhibition in learning allowed and learning prevented conditions of collaborative retrieval.

In Figure 6, we found that collaborative inhibition was still present in the case where models' memories did not homogenize. However, the overall proportion recalled increased in the learning prevented condition and the effect size of collaboration was diminished. We believe that the increase in proportion recalled is a natural effect of how learning works in SAM. 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.

To determine whether this explanation had merit, we checked the average number of words recalled from a list of 50 unrelated words between the learning allowed and learning prevented conditions and found that groups of models in the learning prevented condition recalled significantly more words on average than the groups in the learning allowed condition $t(398) = 22.36$, $p < 0.001$ (see Table 2). This result supports our explanation for why the overall proportion recalled is higher in the learning prevented condition than in the learning allowed condition.

Table 2. Average number of words recalled over 200 simulations of collaborative retrieval

Learning Condition	Mean	SD
Learning Allowed	26.22	3.13
Learning Prevented	34.29	3.42

Additionally, while investigating the effect of each learning parameter on collaborative inhibition in the no learning condition, we found that the collaborative model's *j* parameter made a significant difference in the size of the collaborative inhibition effect. The *j* parameter is a new parameter that was added to the collaborative model to represent learning of the combined group response. This parameter is the incrementing parameter for when word associations are updated within the group response. That is, when a new word is added to the group response, the associations between the new word and the words already stored in the group response are updated according to the *j* parameter. Figure 7 shows the effect of turning the *j* parameter on and off while keeping the other learning parameters (*e, f,* and *g)* off.

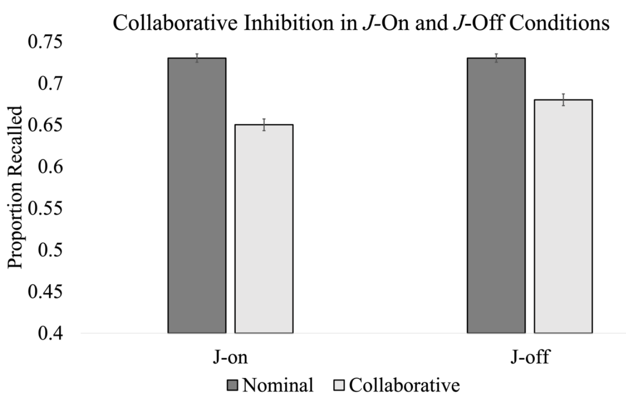

Figure 7. Comparison of collaborative inhibition when the *j* parameter is off and on. All other learning parameters are off in both conditions.

We found that the proportion recalled by the collaborative groups was significantly smaller in the j-on condition than in the j-off condition, $t\ (398) = -4.75$, $p < 0.001$. This means that the *j* parameter may be a significant factor in producing the collaborative inhibition effect.

Discussion

The main conclusion from this study is that memory homogenization during retrieval is not the sole cause of collaborative inhibition in the collaborative SAM model. The goal of part 1 of this study was to determine how models' memories change during collaborative retrieval. We found that models' memories naturally homogenize in a collaborative setting (see Figure 4). When model memories homogenize, they are also less diverse. The goal of part 2 was to determine whether models' memories homogenizing during retrieval is responsible for collaborative inhibition. We found that preventing learning during retrieval also prevented memory homogenization and when learning was prevented, collaborative inhibition persisted, though with a diminished size of effect (see Figure 5 and Figure 6).

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) showed 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. The idea here is 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 shared memories and that group members' memories homogenize due to collaborative recall. The collaborative SAM model shows the pattern of memory homogenization over retrieval (Figure 4) and supports these experimental findings. However, we found that memory homogenization, while present during collaborative retrieval, does not eliminate collaborative inhibition when prevented—suggesting another mechanism, like retrieval disruption, may be involved.

The retrieval disruption hypothesis has the most supporting evidence in the experimental literature (Andersson, Hitch, & Meudell, 2006; Basden, Basden, & Henry, 2000; Finlay, Hitch, & Meudell, 2000) and so we would expect a cognitive model of collaborative memory to support this hypothesis as well. Previously, we showed that the collaborative SAM model supports the retrieval disruption hypothesis (Mannering et al., 2022) when fitting the model to the data from the Basden et al. (1997) study. However, the results of the current study suggest that both retrieval disruption and memory homogenization may play a role in producing collaborative inhibition in the collaborative SAM model.

References

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.

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

Blumen, H. M., & Rajaram, S. (2008). Influence of re-exposure and retrieval disruption during group collaboration on later individual recall. *Memory, 16*(3), 231-244.

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.

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.

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

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

Mannering, W. M., Rajaram, S., Shiffrin, R. M., & Jones, M. N. (2022). Towards a Cognitive Model of Collaborative Memory. *Manuscript in preparation*.

Marion, S. B., & Thorley, C. (2016). A meta-analytic review of collaborative inhibition and postcollaborative memory: Testing the predictions of the retrieval strategy disruption hypothesis. *Psychological Bulletin, 142*(11), 1141-1164.

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

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

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.