

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Why Two Heads Together are Worse Than Apart: A Context-Based Account of Collaborative Inhibition in Memory Search

Permalink

<https://escholarship.org/uc/item/03g2m2rs>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Angne, Hemali

Cornell, Charlotte

Zhang, Qiong

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Why Two Heads Together are Worse Than Apart: A Context-Based Account of Collaborative Inhibition in Memory Search

Hemali Angne^a, Charlotte A. Cornell^b, Qiong Zhang^{a,b,c}

{h.angne, charlotte.cornell, qiong.z}@rutgers.edu

^aDepartment of Computer Science, ^bDepartment of Psychology, ^cRutgers Center for Cognitive Science
Rutgers University–New Brunswick, Piscataway, NJ, USA 08854

Abstract

Contrary to common intuition, groups of people recalling information together remember less than the same number of individuals recalling alone (i.e., the collaborative inhibition effect). To understand this effect in a free recall task, we build a computational model of collaborative recall in groups, extended from the Context Maintenance and Retrieval (CMR) model which captures how individuals recall information alone (Polyn, Norman, & Kahana, 2009). We propose that in collaborative recall, one not only uses their previous recall as an internal retrieval cue, but also listens to someone else's recall and uses it as an external retrieval cue. Attending to this cue updates the listener's context to be more similar to the context of someone else's recall. Over an existing dataset of individual and collaborative recall in small and large groups (Gates, Suchow, & Griffiths, 2022), we show that our model successfully captures the difference in memory performance between individual recall and collaborative recall across different group sizes from 2 to 16, as well as additional recall patterns such as recency effects and semantic clustering effects. Our model further shows that the contexts of collaborating individuals converge more than the contexts of individuals who recall alone. We discuss the contributions of our modeling results in relation to previous accounts of the collaborative inhibition effect.

Keywords: collaborative inhibition; group recall; memory search; computational modeling

Introduction

In daily life, we often rely on others to help us remember things. Consider, for instance, when recalling scenes from a movie with a friend, listening to their memories can help remind you of more plot details. Yet surprisingly, a collaborative group recalls less information than a nominal group – the same number of individuals recalling alone (the *collaborative inhibition effect*, Weldon & Bellinger, 1997). This effect has been robustly observed across a range of empirical situations (for a review, see Marion & Thorley, 2016), including in-person and online settings, small and large groups (Gates et al., 2022), for multiple memory tasks (free recall, Weldon & Bellinger, 1997; cued recall, Kelley, Reysen, Ahlstrand, & Pentz, 2012; and recognition tasks, Andersson & Rönnerberg, 1996), and with various study material of words (Gates et al., 2022), stories and pictures (Weldon & Bellinger, 1997), and film (Wessel, Zandstra, Hengeveld, & Moulds, 2015).

Why does a collaborative group recall less information than a group recalling individually? Multiple accounts have been proposed to explain this counter-intuitive effect. For one, the account of retrieval disruption theorizes that each

individual has an idiosyncratic retrieval strategy that is disrupted when listening to others' recalls, giving rise to non-optimal recalls (B. H. Basden, Basden, Bryner, & Thomas, 1997; Weldon & Bellinger, 1997). Alternatively, the retrieval inhibition account argues that an additional mechanism operates whereby when a group member recalls an item, as-yet-retrieved items are suppressed in listeners' memory, making the remaining items less likely to be retrieved (Barber, Harris, & Rajaram, 2015). Hyman, Cardwell, and Roy (2013) proposed the account of limited exploration in which members of a collaborative group constrain other members' exploration of memory, as they found that collaborative dyads – consisting of two members recalling together – reached fewer categories than nominal dyads when recalling categorized lists.

How can these multiple accounts be understood together? In the current work, we propose context as a unifying theory across individuals and groups. Past work has recognized the important role of context in the encoding and retrieval of information (Anderson & Bower, 1972; Bower, 1967; Estes, 1955; Murdock, 1997). Computational models of memory search built upon context, such as the Context Maintenance and Retrieval (CMR) Model, have successfully captured various behavioral patterns during a free recall task including primacy effects, recency effects, and contiguity effects (Howard & Kahana, 2002a; Lohnas, Polyn, & Kahana, 2015; Polyn et al., 2009). The goal of our current work is to extend CMR, a context-based model of individual recall, to a model of collaborative recall in a free recall task. We will show that the collaborative inhibition effect naturally emerges from the interaction of individuals' mental contexts as they recall information together. In a free recall task, participants study a list of items and are later asked to recall as many items as possible from the list in any order (Murdock, 1962; Roberts, 1972). In CMR, as an individual searches their memories to recall items one after another, each new recall is driven by the context induced by the last recall. For example, when recalling items on a shopping list, having recalled the item 'apple' will induce a fruit-related context in our mind, so that the next recall is likely to be a related item, like 'pear' or 'banana'. To extend individual recall to collaborative recall, we only introduce one additional process to describe how listening to others' recall affects one's own recall: each new recall is now driven by the context induced by either the last recall from the same individual or the last recall from other individ-

uals, controlled by the probability p_{cue} . In other words, when recalling items on a shopping list with other individuals, listening to someone else recall the item ‘apple’ has a similar effect as recalling the item ‘apple’ oneself in which a fruit-related context is induced in one’s mind and leads to the next recall likely being another fruit-related item.

Formulated this way, we will demonstrate in our model simulations that individuals in the nominal condition search their memories independently; in the collaborative group, as recall unfolds, each member’s recall starts to converge to the contexts of others’ recalls. Intuitively, consider memory search in a context space to be analogous to foraging mushrooms in a forest. When individuals explore their memory (or a forest) alone, only their own retrieval of memories (or mushrooms) guides their search. When people explore collaboratively, they guide each other’s search and constrain where one might otherwise be able to explore. This context-based account also helps unify previous accounts. Given the way recall is affected by others in the collaborative group, the trajectory of items recalled in the context space is altered or disrupted compared to that of the nominal condition, similar to what is described in the retrieval disruption account (B. H. Basden et al., 1997). We mathematically formulate where the disruption comes from as well as simulate its effect on recall performance. Additionally, the way contexts converge in our simulations captures the idea that members in a collaborative group constrain other members’ exploration of memory (Hyman et al., 2013). Compared with the retrieval inhibition account (Barber et al., 2015), we provide a simpler account that captures the effect of inhibition without introducing an additional inhibition mechanism that directly suppresses the memory strengths of items.

To test our hypothesis that a context-based account can explain the collaborative inhibition effect, we build a computational model of individual recall and a model of collaborative recall, and we compare their behavior to an existing dataset of nominal and collaborative recall in small and large groups (Gates et al., 2022, Exp. 2). Critical to our modeling approach, we assume that the model of collaborative recall inherits the same parameter values from individual recall (which we obtain from fitting our model to participants in the nominal condition). To account for memory behavior in the collaborative condition, we only have one free parameter, p_{cue} , which describes the probability of listening to others’ recalls. This way, it is guaranteed that any differences we observe between the nominal condition and the collaborative condition are not a result of differences in how members in the two conditions encode and recall information from their mental context, but from the way contexts interact with each other in a group. The fundamental memory processes of how one searches their memories remain the same across individuals whichever condition they are in.

To foreshadow our results, our model successfully captures the difference in memory performance between the nominal and collaborative conditions – the collaborative inhibition ef-

fect – across different group sizes from 2 to 16. Our model also captures additional recall patterns such as recency effects (enhanced recall of items from the end of the list; Murdock, 1962) and semantic clustering effects (semantically similar items are recalled successively; Howard & Kahana, 2002b).

Method

In this section, we first provide an overview of the Gates et al. (2022) study of collaborative recall. We then review the Context Maintenance and Retrieval model upon which we build our model of collaborative recall (CMR; Polyn et al., 2009; Howard & Kahana, 2002a; Lohanas et al., 2015). CMR was developed to explain behavioral patterns observed in a free recall task. As this is the same task that the nominal condition completed, we simply simulate recall in the nominal condition by using multiple CMR models simultaneously. To model collaborative recall, we propose an extension to the retrieval phase of CMR by introducing a probabilistic external cuing mechanism to capture how group members are affected by each others’ recalls.

Gates et al. (2022)’s Study of Collaborative Recall

Gates et al. (2022) conducted a group recall study through Amazon Mechanical Turk. We use their experiment 2 dataset ($N = 1,076$) with groups ranging from size 2 to 16. Participants were assigned to either a nominal or collaborative condition and were not allowed to repeat the task. For each condition, there were 48 groups of size 2, 32 groups of 3; 24 groups of 4, and 12 groups of 8 and 16.

Their experiment consisted of two phases: study and recall. During the study phase, participants individually viewed 60 uncategorized words. Each group saw a different list and, within a group, the presentation order was randomized for each participant. A 30-second long arithmetic filler task then followed. During the subsequent recall phase, in the nominal condition, participants were placed alone in a chatroom and told to type as many words as they could recall from the list into a textbox. In the collaborative condition, participants within a group were placed in a chatroom together and took turns recalling. Recall proceeded in “rounds” during which each participant, in a randomized order, was given 5 seconds to recall a word from the list by typing it into the chatroom. If a participant recalled a not-yet-submitted word, a computerized voice read it aloud to other participants; otherwise, participants could continue attempting to recall a new list item, wait for time to elapse, select a ‘pass’ option, or select a ‘I can’t recall anymore’ option (which ended their participation). There was no time limit for either condition.

Our Proposed Model: Extending CMR to Collaborative Recall

Study Phase (nominal and collaborative conditions) During the study phase, participants in nominal and collaborative conditions individually studied 60 words. To model this, we follow the exact processes of CMR and assume that simulated members of both nominal and collaborative groups encode

a list of items in the same way. When an individual studies an item from the list, their current context drifts towards the memory representations of the recently encountered item. The state of the context at time t is given by

$$c_t = \rho c_{t-1} + \beta_{enc} c^{IN} \quad (1)$$

where c^{IN} is the retrieved context induced by the encountered item, parameter $\beta_{enc} \in [0, 1]$ determines the rate at which context drifts toward that presented item's context c^{IN} , and ρ is a normalizing scalar that renders $\|c_t\| = 1$. The presented list item activates its pre-experimental context c^{IN} :

$$c^{IN} = M_{pre}^{FC} f_t \quad (2)$$

where M_{pre}^{FC} stores item-to-context associations that existed prior to the experiment, and f_t is a binary vector that is all zeros except at the presented item's position. Therefore, $M_{pre}^{FC} f_t$ is the context previously associated with the presented item. In addition to these fixed pre-experimental item-to-context associations held in M_{pre}^{FC} , there are also experimental item-to-context and context-to-item associations held in M_{exp}^{FC} and M_{exp}^{CF} that capture new learning in the experiment. These matrices are initialized to zero and are updated during the study phase. Specifically, when an item is presented, a new association is formed via the Hebbian outer-product learning rule:

$$\Delta M_{exp}^{FC} = \Delta M_{exp}^{CF} = f_t c_{t-1}^T \quad (3)$$

The overall effect of having context drift slowly towards each newly presented item, together with the process of associative learning, is that items presented nearby in the study list tend to be associated with similar context states.

Recall Phase (nominal condition) During the recall phase, participants in the nominal condition recalled items separately while participants in the collaborative condition recalled items together. The recall process for individuals in the nominal condition follows the same recall process as the CMR model. Specifically, during recall, the current context $c_{t-1,j}$ of a simulated participant j drifts towards the retrieved context of the just recalled item c_{rec}^{IN} :

$$c_{t,j} = \rho c_{t-1,j} + \beta_{rec} c_{rec}^{IN} \quad (4)$$

Here, context continues to drift during recall following the same process during study (Equation 1), but at a different rate as determined by $\beta_{rec} \in [0, 1]$ and with c^{IN} expressed differently. During study, an item only retrieves its pre-experimental context when it is presented; however, when an individual recalls an item, its retrieved context activates both its pre-experimental context ($M_{pre}^{FC} f_t$) and its experimental context formed during study ($M_{exp}^{FC} f_t$). The extent of retrieving an item's pre-experimental versus experimental context is determined by a parameter, $\gamma_{fc} \in [0, 1]$, such that:

$$c_{rec}^{IN} = (1 - \gamma_{fc}) M_{pre}^{FC} f_t + \gamma_{fc} M_{exp}^{FC} f_t \quad (5)$$

Once context drifts towards this retrieved context, which items are likely to be recalled? The support (or activation) a^{IN} at time t for recalling different items depends on both how much the current context c_t matches with items' experimental contexts (stored in M_{exp}^{CF}) as well as items' pre-experimental contexts (stored in M_{pre}^{CF} ; see description below). The relative activation of these associations is determined by a parameter, $\gamma_{cf} \in [0, 1]$, such that:

$$a^{IN} = \gamma_{cf} \phi_t M_{exp}^{CF} c_t + (1 - \gamma_{cf}) M_{pre}^{CF} c_t \quad (6)$$

Here, M_{pre}^{CF} is a matrix representing pre-experimental context-to-item associations. To capture semantic clustering effects observed at recall, M_{pre}^{CF} begins as an identity matrix and each element in M_{pre}^{CF} , with indices m and n , is additionally incremented by a semantic association between items m and n , determined by taking the cosine similarity of the two items' GloVe model embeddings (Pennington, Socher, & Manning, 2014). Each entry of the semantic association is additionally raised to the power of λ before being scaled by the constant s_{cf} to match with human semantic representations.

To simulate which item to recall next based on items' support in a^{IN} , the model also needs a retrieval rule and a stopping rule. We use the softmax function as the retrieval rule, $p_i = e^{ka_i^{IN}} / \sum_j e^{ka_j^{IN}}$, where a_i^{IN} is the support to retrieve item i and the parameter k determines the amount of noise. Once an item is retrieved, the context state drifts again, towards that item's retrieved context following Equation (5); cuing with this updated context state supports the retrieval of new items. This retrieval and context updating processes continue until what is determined by a stopping rule: the probability of stopping at each time point is expressed as $p_{stop} = e^{-\epsilon_d a_{nr}^{IN} / a_r^{IN}}$, where a_r^{IN} indicates the summed support for already-recalled items, a_{nr}^{IN} indicates the summed support for not-yet-recalled items, and ϵ_d is a scaling factor (Cornell, Norman, Griffiths, & Zhang, 2024; Kragel, Morton, & Polyn, 2015; Zhang, Griffiths, & Norman, 2023). Overall, because items studied nearby in the list are tied to similar context states during encoding, subsequent recalls are likely to be nearby items in the list; they are additionally likely to be items semantically similar to the current context.

Recall Phase (collaborative condition) While individuals in the nominal condition update their context using only their own recalls, individuals in the collaborative condition additionally can listen to recalls of other members in their group. In Gates et al. (2022), collaborative recall proceeds in "rounds" in which each group member was randomly selected and given a chance to recall. Thus for every recall round in our model, we randomly selected a simulated participant to recall an item. When a simulated participant j recalls an item in our model, their own context drifts towards that item's retrieved context following Equation (5). Because everyone else in the chatroom could hear this recalled item, we let all other simulated participants $i \neq j$ have the chance to use this item as a retrieval cue under probability $p_{cue} \in [0, 1]$

by drifting their internal context $c_{t-1,i}$ towards the cue’s retrieved context $c_{cue,i}$. Under probability $1 - p_{cue}$, participants ignored this item and maintained their current context, $c_{t-1,i}$:

$$c_{t,i} = \begin{cases} p c_{t-1,i} + \beta_{rec} c_{cue,i}, & w.p. p_{cue} \\ c_{t-1,i}, & w.p. (1 - p_{cue}) \end{cases} \quad (7)$$

Notice that the same parameter β_{rec} governs how much one’s context drifts towards someone else’s recall compared with that towards one’s own recall in Equation (5). Intuitively, consider the shopping list example. The context of an individual in the nominal condition updates only according to their own recalls – recalling ‘apple’ from the shopping list induces a fruit-related context so that their next recall will likely be another fruit-related item. In contrast, a collaborative group member’s context updates either according to their own recall or the recall of someone else in their group – recalling ‘apple’ but then listening to another group member’s recall of ‘stapler’ induces a context related to office supplies leading the next recall to likely be an office item like ‘paper’.

For the nominal condition model parameters, we used Bayesian optimization (Mockus, 1998) to search the space of parameters to minimize the normalized root-mean-square error between our model simulations and the nominal condition data. For the collaborative condition model parameters, critically, we assume that the model of collaborative condition inherits the same parameter values from the nominal condition, as the fundamental memory processes (i.e., items’ organization in the context space, the retrieval rule, the stopping rule) stay the same across participants in both conditions. The collaborative model additionally has one parameter, p_{cue} , describing the probability of listening to others’ recall.

Results

Nominal Condition Recall Behavior To simulate the experiment in Gates et al. (2022) with our proposed model, we first fit our model parameters to the free recall behavior in the nominal condition across three sets of behavioral patterns: (1) how well on average words are retrieved for each position in the study list (serial position curve; Murdock, 1962), (2) where in the serial position recall is initiated from (probability of the first recall; Murdock, 1962), (3) how likely it is to recall semantically similar words at adjacent (lag 1) versus far-apart recall positions (lag 2, 3, 4). Semantic similarity is computed by finding the average cosine similarity between every pair of recalled items at different lags for their output positions (Cornell et al., 2024).

Figure 1 compares the observed patterns in the data and the model (nRMSE = -6.80). The serial position curve (Figure 1A) and probability of first recall (Figure 1B) in the nominal condition displayed recency effects – enhanced recall of items from the end of the list (Murdock, 1962). The model was able to capture recency effects because one’s context at the start of recall is most similar to the contexts of the last few studied items. There are also semantic clustering effects in the nominal condition (Figure 1C), as we observed higher semantic similarity at a small lag 1 than at a large lag 4 in both




Figure 1: Behavioral recall patterns of individuals in the *nominal condition* and the model fit. These patterns are (A) serial position curve, (B) first recall probability, and (C) semantic similarity by lag. The shaded error represents the standard error of the mean. The parameter set was: $\beta_{enc} = 0.594$, $\beta_{rec} = 0.871$, $\gamma_{fc} = 0.297$, $\gamma_{cf} = 0.344$, $s_{cf} = 0.685$, $k = 2.693$, $\epsilon_d = 0.585$, $\lambda = 1.751$.

participant data ($t(8803) = 5.585$, $p < .001$) and the model ($t(49227) = 6.074$, $p < .001$). The model was able to capture semantic clustering effects because retrieving an item ‘apple’ updates the current context to make it easier to retrieve a semantically related item ‘pear’. Note that some typical free recall behaviors were not observed in the data such as primacy and temporal contiguity effects, potentially due to the long list length of sixty in the human experiment. For the rest of the analyses, we focus on modeling behavioral patterns of the observed recency effects, semantic clustering effects, and memory performance.

Collaborative Condition Recall Behavior Our model could capture the recall behavior of individuals in the nominal condition. We next tested if the same set of parameters could also capture the recall behavior of individuals in the collaborative condition. The collaborative condition inherits its parameter set from the nominal condition as we hypothesized that the fundamental memory search processes (i.e., how context is used to encode and later retrieve items from the context space) are the same in nominal and collaborative conditions. The collaborative condition additionally has only one parameter, p_{cue} , describing the probability of listening to others’ recall. To fit p_{cue} , we searched for a value from 0 to 1 in 0.1 increments that minimized the normalized root-mean-square error between our model simulations and the collaborative condition data. The fitting was done across the same behavioral patterns as in the nominal condition (Figure 2A-C) as well as an additional behavioral pattern (Figure 2D) that characterizes the process of listening to the recalls of others. While Figure 2C illustrates how a recall is semantically related to previous recalls of the same individual (as participants in the collaborative condition may attend to their own recalls), Figure 2D illustrates how a recall is semantically related to previous recalls of the entire group (as participants may attend to others’ recall in addition to their own recalls).

Participants in the collaborative condition showed typical




Figure 2: Behavioral recall patterns for individuals in the *collaborative condition* and the model fit. These patterns are (A) the serial position curve, (B) the probability of first recall, and semantic similarity by lag related to (C) previous recalls of the same individual versus (D) the entire group. The shaded error represents the standard error of the mean. $p_{cue} = 0.2$, with the rest of the model parameters inherited from the nominal condition.

free recall behaviors similar to that of the nominal condition, demonstrating recency effects (Figure 2A-B) and semantic clustering effects (Figure 2C). Participants not only tend to retrieve items that are semantically related to their own recall (Figure 2C), with higher semantic similarity at lag 1 than at lag 4 (participant data: $t(4235) = 4.009$, $p < .001$; model: $t(17307) = 3.292$, $p < .001$), but also retrieve items that are semantically related to others' recall (Figure 2D; participant data: $t(5504) = 2.143$, $p = .03$; model: $t(28582) = 2.680$, $p = .007$). Our model was able to capture these patterns with its parameters fit to the recall behavior of individuals in the nominal condition. These results support our hypothesis that one not only uses their previous recall to drive their recall context but also listens to someone else's recall.

Next, we test the key hypothesis of the present work and examine if our context-based model can capture the collaborative inhibition effect. We examined this under multiple values of p_{cue} , instead of just the best-fit value, to determine if the collaborative inhibition effect can emerge as long as there is some context interaction between individuals ($p_{cue} > 0$). Figure 3A shows the effect of group size on the amount of collaborative inhibition in Gates et al. (2022), whereby as group size increases, inhibition first increases and then decreases. Our model was able to capture this qualitative trend under different values of $p_{cue} > 0$ (Figure 3B). These results support that the performance difference between the collaborative and the nominal condition arose naturally from our model's collaborative mechanism and was not sensitive to the values of the parameter we introduced to model collaborative recall.




Figure 3: The context-based model of collaborative recall captures the increase and subsequent decrease in inhibition by group size. (A) Collaborative inhibition observed in Gates et al. (2022), computed as performance difference between the nominal condition and the collaborative condition. (B) Collaborative inhibition predicted by the model under different values of p_{cue} . The shaded error represents the standard error of the mean.

Context Dynamics Our model captured the collaborative inhibition effect at different group sizes. Why and how does the collaborative inhibition effect arise in the model? We hypothesized that when individuals interact in a collaborative setting, they listen to each other's recall, and their mental contexts become synchronized over time. Compared with the nominal condition where individuals use diverse and unique contexts, synchronized contexts in the collaborative group may constrain one's ability to recall. To evaluate if listening more to others' recall (increasing p_{cue}) gives rise to more synchronized contexts in a group, we measured context similarity in a group by computing the average cosine similarity between context vectors of all possible pairs of simulated members per group after each recall. Figure 4 plots the context similarity across recall outputs during the recall phase for both nominal and collaborative groups of sizes 3 and 16. Consistent with our hypothesis, our model simulations show that




Figure 4: The model shows increasing context similarity in collaborative groups during recall. Average cosine similarity between simulated group members' context states during recall for parameter values $p_{cue} > 0$ increased more in the collaborative groups than in the nominal condition for group sizes (A) 3, and (B) 16. The shaded error represents the standard error of the mean.

participants in the collaborative condition increased context similarity more quickly and maintained greater context similarity than in the nominal condition. We performed a two-way ANOVA to examine if the change in context similarity from early to late recalls (0, 10) is different across nominal and collaborative ($p_{cue} = 0.3$) conditions, and there was a significant interaction between recall condition and recall output for group size 3 ($F(1, 396) = 193.689, p < .001$) and group size 16 ($F(1, 396) = 1052.515, p < .001$). Similarly, we also observed greater context similarity with larger values of p_{cue} , as there was a significant interaction between different collaborative conditions ($p_{cue} = 0.3, p_{cue} = 0.7, p_{cue} = 1.0$) and recall output for group size 3 ($F(1, 1196) = 584.980, p < .001$) and group size 16 ($F(1, 1196) = 3178.810, p < .001$). When $p_{cue} = 0$, context similarity in nominal and collaborative conditions is similar as collaborative group members are no longer affected by others' recalls. These results support that listening to each others' recalls in a collaborative setting guides individuals to reach more similar areas of the memory space, giving rise to more synchronized mental contexts.

General Discussion

Multiple accounts have been proposed to explain why collaborative groups recall less information than nominal groups (Barber et al., 2015; B. H. Basden et al., 1997; Hyman et al., 2013). We proposed a context-based account to unify these accounts; specifically, in collaborative recall, the context of an individual not only evolves on its own but is also influenced by the context of others in the group. To test our account, we built a computational model of collaborative recall in groups, extended from the Context Maintenance and Retrieval (CMR) model which captures how individuals recall information alone (Polyn et al., 2009). By comparing our model's simulations to an empirical dataset of nominal and collaborative group recall (Gates et al., 2022, Exp. 2), we found it is able to capture the collaborative inhibition effect across different group sizes as well as the recency effects and the semantic similarity effects. We also found that collaborative group members' context convergences in the context space more than in the nominal condition. These results support our proposed account of collaborative inhibition: Minds within a collaborative group become aligned or synchronized with each other, thus missing opportunities to recall unique information that others may not have considered.

Our proposed context-based model could be extended in the future to account for a related paradigm in external cuing literature, namely – part-set cuing (Slamecka, 1968), given the similarity between a part-set cuing paradigm and a collaborative recall paradigm. In collaborating recall, external retrieval cues are provided by the recalls of other group members; in part-set cuing, individuals complete a free recall task in which, after studying a list of items, some participants are provided a subset of list items by the experimenters as retrieval cues. These cued participants recall fewer of the non-cue items than participants who did not receive any cues (for

reviews, see Bäuml, 2007; Nickerson, 1984). Part-set cuing has also been explained by similar accounts as collaborative inhibition, namely retrieval disruption (D. R. Basden, Basden, & Galloway, 1977) as well as retrieval inhibition (Bäuml & Aslan, 2004; Bäuml & Aslan, 2006). However, context may also account for the reduced recall performance observed in cued individuals. Following a context-based account, an item (whether externally provided or internally generated) updates the location of one's current context to be more like the item. Cues may disrupt one's context, making it more similar to the context of cue items than non-cue items. This, in turn, may limit one's ability to recall non-cue items. This would provide a more parsimonious account by not assuming additional mechanisms such as inhibition (for a related context-based account for the effect of a single, experimenter-provided cue, see Cornell et al., 2024).

There have been two other computational models of collaborative recall. One is Luhmann and Rajaram (2015)'s agent-based model in which when one group member recalls an item, it reduces the activations of other items based on their similarity to that recalled item (i.e., following the retrieval inhibition account). A second model extended the Search of Associative Memory (SAM) model (Raaijmakers & Shiffrin, 1981) and proposed that group members listen to the recall of whoever recalls first, and this item is then used to cue their next recall (Mannering, Rajaram, & Jones, 2021). Although this model does not include an inhibitory mechanism, some of the parameters are fit separately for the nominal condition and collaborative condition, which could directly explain the difference in memory performance between the two conditions (i.e., the collaborative inhibition effect). In our proposed model, we do not include inhibition, and we assume that collaborative recall inherits the same memory processes (model parameters) from individual recall in the nominal condition. Therefore, the collaborative inhibition effect captured by the model is not a result of how participants in two conditions encode and recall information differently, but a result of the interaction among participants' mental contexts in the collaborative group.

Our computational modeling work offers unique contributions to collaborative memory research. We provide a strong test of our proposed context-based account by demonstrating its ability to capture key recall patterns in collaborative recall without directly fitting parameters to the participant data in the collaborative condition. We simulate the collaborative group behavior based on model parameters obtained from the participant data in the nominal condition. Our model of collaborative recall is able to capture the collaborative inhibition effect, as listening to others' recall constrains where one could otherwise be able to search in the context space. Moreover, our model can account for detailed patterns of collaborative recall behavior including the recency effects and the semantic clustering effects. Taken together, our results support the important role of context in explaining a range of memory findings across individuals and groups.

References

- Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological Review*, 79(2), 97–123.
- Andersson, J., & Rönnerberg, J. (1996). Collaboration and memory: Effects of dyadic retrieval on different memory tasks. *Applied Cognitive Psychology*, 10(2), 171–181.
- 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.
- Basden, B. H., Basden, D. R., Bryner, S., & Thomas, R. L. I. (1997). A comparison of group and individual remembering: Does collaboration disrupt retrieval strategies? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(5), 1176–1189.
- Basden, D. R., Basden, B. H., & Galloway, B. C. (1977). Inhibition with part-list cuing: Some tests of the item strength hypothesis. *Journal of Experimental Psychology: Human Learning and Memory*, 3(1), 100–108.
- Bäuml, K.-H. T. (2007). Making memories unavailable: The inhibitory power of retrieval. *Zeitschrift für Psychologie/Journal of Psychology*, 215(1), 4–11.
- Bäuml, K.-H. T., & Aslan, A. (2004). Part-list cuing as instructed retrieval inhibition. *Memory & Cognition*, 32(4), 610–617.
- Bäuml, K.-H. T., & 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), 33–43.
- Bower, G. (1967). A multicomponent theory of the memory trace. *Psychology of Learning and Motivation*, 1, 229–325.
- Cornell, C. A., Norman, K. A., Griffiths, T. L., & Zhang, Q. (2024). Improving memory search through model-based cue selection. *Psychological Science*, 35(1), 55–71.
- Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, 62(3), 145–154.
- Gates, V., Suchow, J. W., & Griffiths, T. L. (2022). Memory transmission in small groups and large networks: An empirical study. *Psychonomic Bulletin & Review*, 1–8.
- Howard, M. W., & Kahana, M. J. (2002a). A distributed representation of temporal context. *Journal of mathematical psychology*, 46(3), 269–299.
- Howard, M. W., & Kahana, M. J. (2002b). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, 46(1), 85–98.
- Hyman, J., Ira E, Cardwell, B. A., & Roy, R. A. (2013). Multiple causes of collaborative inhibition in memory for categorised word lists. *Memory*, 21(7), 875–890.
- 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.
- Kragel, J. E., Morton, N. W., & Polyn, S. M. (2015). Neural activity in the medial temporal lobe reveals the fidelity of mental time travel. *Journal of Neuroscience*, 35(7), 2914–2926.
- Lohnas, L. J., Polyn, S. M., & Kahana, M. J. (2015). Expanding the scope of memory search: Modeling intralist and interlist effects in free recall. *Psychological Review*, 122(2), 337–363.
- 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., Rajaram, S., & Jones, M. N. (2021). Towards a cognitive model of collaborative memory. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 43, pp. 959–965).
- 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.
- Mockus, J. (1998). The application of bayesian methods for seeking the extremum. *Towards global optimization*, 2, 117–129.
- Murdock, B. B. (1962, Nov). The serial position effect of free recall. *Journal of Experimental Psychology*, 64(5), 482–488.
- Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory (TODAM2). *Psychological Review*, 104(4), 839–862.
- Nickerson, R. S. (1984). Retrieval inhibition from part-set cuing: A persisting enigma in memory research. *Memory & Cognition*, 12(6), 531–552.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532–1543).
- Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, 116(1), 129–156.
- Raaijmakers, J. G., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological review*, 88(2), 93–134.
- Roberts, W. A. (1972). Free recall of word lists varying in length and rate of presentation: A test of total-time hypotheses. *Journal of Experimental Psychology*, 92(3), 365–372.
- Slamecka, N. J. (1968). An examination of trace storage in free recall. *Journal of experimental psychology*, 76(4p1), 504–513.
- Weldon, M. S., & Bellinger, K. D. (1997). Collective memory: Collaborative and individual processes in remembering. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(5), 1160–1175.
- Wessel, I., Zandstra, A. R. E., Hengeveld, H. M., & Moulds, M. L. (2015). Collaborative recall of details of an emotional film. *Memory*, 23(3), 437–444.

Zhang, Q., Griffiths, T. L., & Norman, K. A. (2023). Optimal policies for free recall. *Psychological Review*, 130(4), 1104–1124.