UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

A LITMUS TEST FOR THE PATTERN-SUPPRESSION MODEL

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

THOMAS MICHAEL BOBBITT
Norman, Oklahoma
2016

A LITMUS TEST FOR THE PATTERN-SUPPRESSION MODEL

A THESIS APPROVED FOR THE
PSYCHOLOGY DEPARTMENT

BY

Dr. Daniel Kimball, Chair

Dr. Scott Gronlund

Dr. Luz-Eugenia Cox-Fuenzalida

This thesis is dedicated to Sarah Elaine Bobbitt.

Acknowledgements

I would like to thank Dr. Daniel Kimball for all of the time and attention he has given to my various research projects, including those reported in this thesis. He has also always encouraged independent thinking on my part, and acted as a guide by informing me of various relevant research and statistical methods.

I would also like to thank my committee members, as well as my fellow members of the Cogfog Lab, for the insight and perspectives that they have provided throughout my project, as well as encouragement and reassurance.

Lastly, I would like to thank my family, who have been so supportive of me and my education. I would like to thank in particular my wife, Sarah.

Table of Contents

List of Tables

List of Figures

Abstract

The current set of experiments were designed to test the *pattern-suppression model*, which is a component of the suppression theory of forgetting. Experiment 1 tested the *pattern-facilitation hypothesis*, a hypothesis derived from the model and which states that the processing of an item will result in increased accessibility of other items from memory that share semantic features with the item. There was no support for this hypothesis. Experiments 2-3 examined the nature of *second-order inhibition*, which is the finding that the direct suppression of an item will result in indirect suppression of other items that share semantic features with the item in question. According to the *pattern-suppression model*, items affected by *second-order inhibition* are only suppressed with respect to the shared semantic features. However, Experiments 2-3 found no evidence for this assumption. In fact, there appeared to be, if anything, facilitation of these features. The results are elaborated upon with respect to the inhibition literature.

A Litmus Test for the Pattern-Suppression Model

The inhibition theory of forgetting is a prominent theory in memory research (for reviews, see Anderson, 2003; Storm & Levy, 2012), but has not been without its detractors (Jonker, Seli, & Macleod, 2013; Raaijmakers & Jakab, 2013). A guiding theme for the theory is that the process of forgetting facilitates our ability to remember and hence is adaptive (Bjork, 1989). This adaptive forgetting has been studied extensively in the context of the very act of remembering: in order to efficiently locate and retrieve the memory that one is trying to remember, one must forget—or suppress—other memories that would otherwise "spring to mind" or "get in the way." A useful way to grasp the necessity of suppression mechanisms is to consider situations in which they fail, such as when one experiences having a memory on the "tip of the tongue" (TOT). Brown and McNeill (1966) relate one such incident:

> For several months we watched for TOT states in ourselves. Unable to recall the name of the street on which a relative lives, one of us thought of *Congress* and *Corinth* and *Concord* and then looked up the address and learned that it was *Cornish*.

In this case, the researcher could not remember *Cornish*; however, the fact that the names he thought of were so similar in sound attests to the fact that *Cornish* was stored in memory. According to the inhibition theory, the researcher simply was not able to successfully suppress the memories of the similar words, which, consequently, impaired his ability to remember the name of the street (for a discussion of suppression failure, see Anderson & Levy, 2011). In the usual situation, one is using a cue (e.g., a physics equation) to remember a particular item that is associated with the cue (e.g., the equation for rotational acceleration) and, according to the theory, suppresses other memories that are strongly associated to the same cue (e.g., other physics equations). According to the inhibition theory, we would have TOT experiences (and other similar failures of memory) all of the time were it not for our ability to suppress memories (Anderson, 2003). The purpose of the current study is to test an important (but too little studied) component of the inhibition theory, namely, the *pattern-suppression model*, which will now be described.

The *pattern-suppression model* is not an original component of the inhibition theory; rather, it was postulated to make sense of Anderson and Spellman's (1995) results. In the first part of their experiment (Experiment 1, see Figure 1), participants (Ps) studied 24 category-exemplar pairs (e.g., *RED-blood*), with 6 exemplars being studied for each of 4 categories. Next, Ps had their memory for 3 exemplars from two of the categories tested on a retrieval-practice test, where each test item consisted of the category cue and the first two letters of the exemplar (e.g., *RED-bl*__). Categories that were used as cues in this phase were termed retrieval-practice (Rp) categories and the tested exemplars were termed Rp+ items; exemplars belonging to the Rp categories that were not tested were termed Rp- items and exemplars belonging to the two other categories were termed Nrp items. According to the inhibition theory, the Ps should have used suppression in order to facilitate their ability to remember the Rp+ items (e.g., *RED-blood*) on this retrieval-practice test. Specifically, they should have suppressed memories of other items associated to the cue, which in this case would include the other exemplars belonging to the category in question (e.g., *RED-apple*). On a final memory test, Ps did have impaired success in remembering Rp- items, which was an original prediction of the inhibition theory. However, Anderson and Spellman (1995) also found impairment for exemplars that neither explicitly nor implicitly belonged to an Rp category, but that did implicitly share a category with other suppressed exemplars (e.g., *FOOD-cracker*. Crackers are not red and hence a memory representation of them does not need to be suppressed when trying to remember a red item; however, crackers are a food, as are apples, the memory representation of which was suppressed). This result was not expected, and *the pattern-suppression assumption* was developed to account for it.

The *pattern-suppression model* states that memory representations are distributed in nature rather than unitary (Anderson & Spellman, 1995). For example, the representation of *apple* is not a single "trace," but rather a grouping of features, such as *red, fruit, round,* etc. Successful retrieval of a memory representation requires that each of its features be activated. Additionally, if two objects (e.g., apple and cracker) share a quality (e.g., they are both foods), then memory representations of those objects will *not* each independently represent that quality; rather, the quality will be represented by a single feature (rather than two) that is *shared* between the two representations (see Figure 2). Assuming that memory representations are distributive has implications for the nature of suppression. If the suppression of a memory entails the suppression of *all* of its features, then some features of the very memory one is trying to remember may themselves be suppressed. This maladaptive effect would occur to the extent that the memory one is trying to remember shares features with the memories one is suppressing. Hence, in order to avoid the maladaptive effect, the *pattern-suppression assumption* further states that the direct suppression of a memory only affects features that are *not* shared with the memory one is trying to remember (Anderson & Spellman, 1995; Goodmon & Anderson, 2011). However, other memories that are not directly suppressed will nonetheless be indirectly affected by suppression if they share features with a memory that is directly suppressed. The reason that Anderson and Spellman (1995) found impairment for items (e.g., *FOOD-crackers*) not belonging to an Rp category (e.g., *RED*) is that those items shared features (e.g., *FOOD* features) with suppressed items (e.g., *RED-apple*). This indirect form of suppression is known as *second-order inhibition.*

Although the *pattern-suppression model* is useful for explaining *second-order inhibition* (as well as *integration effects*; see Anderson, Green, & McCulloch, 2000; Goodmon & Anderson, 2011), there are some results that appear to be inconsistent with it. For example, Anderson, Bjork, and Bjork (2000) showed that when subjects merely *restudy* a subset of exemplars before the final test (rather than retrieve the exemplars on a practice test) there is no impairment for the other exemplars belonging to the same categories as the restudied exemplars. This finding is said to be consistent with the inhibition theory because suppression should only occur when one is attempting retrieval—the purpose of suppression is to facilitate retrieval. However, Raaijmakers and Jakab (2013) have pointed out that, according to the *pattern-suppression model*, these non-restudied exemplars (e.g., *FRUIT-grapes)* do share features with the restudied exemplars (e.g., *FRUIT-orange*) at least insofar as the two groups share a category (*FRUIT*). Furthermore, if restudy of exemplars makes them more accessible during the final test, as it did for Anderson et al. (2000), and if this increased accessibility of the exemplars indicates increased activation of the features belonging to them, then other exemplars that share features with the restudied exemplars should also have an increased accessibility. This will be termed the *pattern facilitation hypothesis*. However, Anderson et al. (2000) found no increased accessibility of non-restudied exemplars even when they shared a category with the restudied exemplars (see also Anderson & Bell, 2001; Bäuml, 2002). One possible explanation is that the increased accessibility was too slight to be detected. Perhaps the feature overlap between the restudied exemplars and the non-restudied exemplars was too minimal to make a detectable difference.

The current study was undertaken with two goals in mind. First, I wished to test the *pattern facilitation hypothesis*. In order to increase its detectability, I have used Goodmon and Anderson's (2011) stimuli. These stimuli consist of category-exemplar pairs, just as did the prior studies that have been discussed. However, the exemplars were selected such that each (e.g., *ANIMAL-lion*) had a very high degree of overlap with one and only one other exemplar within the same category (e.g., *ANIMAL-tiger*). If Ps restudy only exemplars that do *not* have high overlap with each other, then each restudied exemplar will therefore have high overlap with one of the non-restudied exemplars from the same categories. According to the *pattern-suppression model*, these none-restudied exemplars should therefore have an increased accessibility because of the fact that they have high feature overlap with restudied exemplars. This prediction was tested in Experiment 1.

The second goal, worked for in Experiments 2-3, was to re-examine a previously tested prediction of the *pattern-suppression model:* second-order inhibition. However, unlike prior studies, I wished to test this prediction at the *feature* level. That is, I wanted to see if items affected by second-order inhibition are indirectly suppressed with respect to a specific, predetermined subset of their features (i.e., those shared with a directly suppressed item), but not suppressed with respect to a different predetermined subset of their features (i.e., those not shared with a directly suppressed item). By examining the accessibility of the features (rather than memory representation as a whole), it can be determined whether or not the features really are being affected in the ways predicted by the *pattern-suppression model*.

Experiment 1

The first experiment closely followed Goodmon and Anderson (2011). However, before the final test, Ps not only did retrieval-practice on a subset of the items, but were also give extra exposure for another subset of the items. Predictions of the *pattern-suppression model* were as follows. First, because we used Goodmon and Anderson's (2011) stimuli (in which Rp+ items and Rp- items have high semantic feature-overlap *with each other*), retrieval practice of items within a category should result in neither impairment nor facilitation of other items within that same category. This finding would replicate Goodmon and Anderson (2011) and follows from the logic which states that the Rp- items should have both facilitated features (i.e., the features that they share with an Rp+ item) and suppressed features (i.e., the features that are not shared with an Rp+ item), and so therefore should be subject to neutralizing forces, leaving them to be relatively unaffected in their overall accessibility (see Figure 2a). The second prediction concerned the items that did not receive extra exposure themselves (which will be called Ee- items), but did belong to a category that had exemplars that did receive extra-exposure (which will be called Ee+ items). Ee- items had high semantic feature overlap with Ee+ items. Because extra exposure of an item, unlike retrieval practice, does not result in suppression, the non-shared features of Ee- items should not have been suppressed; however, because extra-exposure of Ee+ items should have facilitated their own accessibility, and therefore facilitated their features, the features that Ee- items shared with Ee+ items should have also been facilitated (see Figure 2b). Thus, Ee- items should have received a boost in accessibility from the extra exposure of Ee+ items.

Method

Participants

Forty-nine undergraduates (36 women) from the University of Oklahoma participated in the current study. They ranged from 18 to 22 years of age (*M* = 18.8). The participants (Ps) were predominantly White (82%), with remaining Ps being American Indian or Alaskan Native (2%), Asian (6%), Black or African American (8%), and Native Hawaiian or other Pacific Islander (2%). Ps received class credit for participating.

Design

The status of each memory item was manipulated within-subjects on 8 levels (Rp+, Rp-, Nrp+, Nrp-, Ee+, Ee-, Nee+, and Nee-). Items were Rp+ if they were given retrieval practice prior to the final test; items that belonged to the same category as the Rp+ items and that were not given retrieval practice were Rp- items; items that were given extra exposure prior to the final test were Ee+ items; items that belonged to the same category as Ee+ items and that were not give extra exposure were Ee- items; the remaining items were baseline items, with Nrp+ items serving as a baseline for Rp+

items, Nrp- items for Rp- items, Nee+ items for Ee+ items, and Nee- items for Ee- items. For the final test, subjects were given category-plus-stem cues (e.g., *ANIMAL-h___*), with all 8 exemplars of a category being tested before moving to the exemplars of the next category. The Rp-, Ee-, Nrp-, and Nee- items were tested in the first four positions of their respective categories (there were 4 categories total); the Rp+, Ee+, Nrp+, and Nee+ items were tested in the last four positions of their respective categories.

Procedure

Ps provided verbal consent and were then informed that they were participating in an experiment on "memory and reasoning." (They were told the experiment covered reasoning so that later distractor tasks, which concerned reasoning, would be believed to be a part of the study.) They were then told that they would be presented with 48 category-exemplar pairs (16 were fillers), and that they should study each pair by relating the exemplar to the category. Ps were then presented with the pairs and had 5 s to study each. The pairs were presented in a randomized block, with each block containing 1 exemplar from each category. Furthermore, in order to minimize primacy and recency effects, the first 3 and last 3 pairs presented were filler items. In order to further rule out confounding factors, the following restrictions were placed on the presentation order: no 2 categories appeared in sequence more than once; successive exemplars were never from the same category; there was an average of 7 intervening pairs in between successive exemplars from a given category; Rp+, Rp-, Ee+, and Ee- items were distributed evenly throughout the study phase. Filler items were used to meet these constraints. Important to note: if a given exemplar (e.g., *ANIMAL-lion*) was an Rp+ item (or Ee+, Nrp+, or Nee+ item), then its semantic pair (e.g., *ANIMAL-tiger*) was an Rp- item (or Ee-, Nrp-, or Nee- item, respectively).

After Ps completed the study phase, they received further instruction for the second phase of the experiment. First, they were told that they would be tested over the pairs; each test item would present a category and the first two letters of one of its exemplars (e.g., *ANIMAL-li*__) and they had 10 s to recall and type the appropriate exemplar. They were then told that, during the test, they would also have an opportunity to restudy some of the items. Specifically, they were told that when a complete word pair (e.g., *PROFESSION-doctor*) appeared, that they had 10 s to type the complete pair as well as to study it.

Ps were given 3 retrieval-practice trials on each of 4 exemplars from 1 category and 3 extra-exposure trials on each of 4 exemplars from another category. There were 3-4 intervening trials (average: 3.75) in between the first and second trials of a given exemplar, and 5-8 intervening trials (average: 6.75) between the second and third trial of a given exemplar. Thus, there was an expanding schedule. Furthermore, exemplars from the same category were never presented on successive trials and there were no repeating sequences. Fillers were used throughout in order to maintain these constraints.

After Ps finished the second phase, they worked on a distractor task for 5 min before taking a final test over all of the presented pairs. Each test presented a category and the first letter of one of its exemplars (e.g., *ANIMAL-l*___), and Ps had 10 s to type in the appropriate word. Ps were tested over every exemplar from a given category before moving on to the next category. After Ps completed the final test, they answered questions concerning study strategies they used during the experiment. Ps were then thanked for their participation.

Materials

Ps studied 8 exemplars for each of 4 unrelated categories (*ANIMAL, PROFESSION, APPLIANCE, WEAPON*), and so studied a total of 32 word-pairs (e.g., *ANIMAL-lion*). Eight exemplars of an additional 2 categories (*CITY, RIVER*) were also studied, but merely served as fillers throughout the experiment. The stimuli used were the same as those used by Goodmon and Anderson (2011, see Appendix A). I chose these stimuli because Goodmon and Anderson (2011) have already shown that they lead to semantic integration effects (i.e., elimination of an RIF effect that is simply due to the semantic natures of the Rp+ and Rp- items), thereby suggesting, on the *pattern-suppression model*, that there is a high amount of semantic overlap between each exemplar (e.g., *lion*) of a category (e.g., *ANIMAL*) with another exemplar (e.g., *tiger*) of that same category.

A few counter-balancing measures were taken in order to rule out confounds. First, the baseline categories (i.e., categories that were not presented during the second phase for additional study or for testing) were tested before the comparison categories (i.e., the Rp and Ee categories) for half of the Ps and were tested after the comparison categories for the other half of the Ps. Second, *ANIMAL* (WEAPON) was the Rp (Nrp) category for half of the Ps and the Nrp (Rp) category for the other half of the Ps. Similarly, *PROFESSION* (*APPLIANCE*) was the Ee (Nee) category for half of the Ps and the Nee (Ee) category for the other half of the Ss. The status of an item within a category (e.g., Rp+ vs. Rp-) was determined randomly, with the following constraint: if a given exemplar (e.g., *ANIMAL-lion*) was an Rp+ item (or Ee+, Nrp+, or Nee+ item), then its semantic pair (e.g., *ANIMAL-tiger*) was an Rp- item (Ee-, Nrp-, or Nee- item, respectively).

For the distractor task, Ps worked on a Morningness—Eveningness Questionnaire (Rifkin, Jacobs, & White, 2001), a Cognitive Failures Questionnaire (Broadbent, Cooper, Fitzgerald, & Parkes, 1982) and causal-reasoning problems for a total of 5 min. At the end of the experiment, Ps were given a questionnaire to complete. An item assessing covert retrieval activity was included, in which Ps were asked how often (scale of 1-6; 1 = none of the time; 6 = all of the time) "When I was asked to remember a specific exemplar for a memory test, I would use that time to also think of other exemplars that I had studied."

Results

In the current experiments, Bayesian analyses were used in addition to traditional *t*-tests and regression analyses. The benefit of a Bayesian analysis is that, unlike traditional statistical tests, it can detect evidence supporting both the alternative hypothesis and the null hypothesis. The output statistic is called the Bayes factor (B_{01}) and it is a ratio of the posterior odds of the null hypothesis being true to the posterior odds of the alternative hypothesis being true. A Bayes factor of at least 3 constitutes moderate evidence in favor of the null hypothesis and a Bayes factor of 10 or above constitutes strong evidence in favor of the null; a Bayes factor of 1/3 or less constitutes moderate evidence in favor of the alternative hypothesis and a Bayes factor of 1/10 or less constitutes strong evidence in favor of the alternative hypothesis. Bayes factors were calculated at http://pcl.missouri.edu/bf-one-sample and the Jeffrey—Zellner—Siow (JZS) prior was used.

Retrieval-practice of Rp+ items led to facilitated performance on the final test for those items (M = .74, SD = .26) in comparison to Nrp+ items (M = .51; SD = .20), t (48) = 5.4, $p < .05$, Cohen's d = .99, B_{01} = 0001; additionally, performance on the final test over Rp- items (M = .44, SD = .27) did show marginal evidence of impairment in comparison to Nrp items (M = .52, SD = .25), t (48) = -1.9, p = .059, Cohen's d = -.31, B_{01} = 1.15. Such impairment would be important because it would contradict the findings of Goodmon and Anderson (2011), who found that Rp- items are not impaired when they have substantial semantic overlap with Rp+ items. However, the Bayes Factor was in favor of the null hypothesis (albeit *very* weakly), suggesting that, if anything, the results are in accord with Goodmon and Anderson (2011). The Bayes Factor value is difficult to interpret, however, for it gives no substantial support to either the impairment interpretation or to the no-impairment interpretation. Thus, further analyses were run.

Specifically, I looked at how well Rp- items (e.g., *lion*) were recalled given either that the overlapping Rp+ pair (*tiger*) was successfully recalled during retrieval-practice the participant or that it was not successfully recalled. If the overlapping Rp+ item was not successfully recalled during the prior retrieval practice, then the Rp- item should suffer from suppression (Storm & Nestojko, 2010), but not receive any facilitation, meaning that impairment *is* predicted for the Rp- item in this case. It is only when the overlapping Rp+ item is successfully recalled during the prior retrieval practice that the Rp- item should be unaffected, according to the *pattern-suppression* *model*. Ps recalled 48% of Rp- items when they had earlier successfully recalled the overlapping Rp+ pair during retrieval practice, but only 9% of Rp- items were recalled when Ps had not successfully recalled the overlapping Rp+ pair during retrieval practice. To compare these two conditional performances with a paired *t*-test, Ps who recalled all of the Rp+ items had to be excluded (they had no percentage of Rp- items recalled given that the overlapping Rp+ items were not recalled), resulting in Ps who merely recalled 36% of Rp- items when they had earlier successfully recalled the Rp+ pair during retrieval practice. Nevertheless, a significant difference was still detected [t (23) = 3.44, $p < .05$ [= .002], Cohen's d = 1.01, B_{01} = .057] and was in the direction predicted by the *pattern-suppression model*. Furthermore, when comparing performance of Rp- items *that overlapped with Rp+ items that were successfully recalled on retrieval-practice test* (M = .48, SD = .29) to Nrp- items (M = .52, SD = .25), no significant difference was found, t (48) = .69, p = .50, Cohen's d = .15, B_{01} = 5.14. Hence, there is really no opposition to the data of Goodmon and Anderson (2011) on this point.

Extra-exposure of Ee+ items led to facilitated performance on the final test for those items (M = .66, SD = .24) in comparison to Nee+ items (M = .56, SD = .26), t (48) = 2.5, $p < .05$, Cohen's d = .40, B_{01} = .4; however, there was not a facilitated performance for Ee- items (M = .59, SD = .23) in comparison to Nee- items (M = .56, SD = .23), t (48) = .69, p = .50, Cohen's d = .13, B_{01} = 5.16, which contradicts the original prediction of the *pattern-suppression model*.

The most important data concerned how well the Ee- items were recalled. In contrast to the prediction of the *pattern-suppression model*, no significant facilitation was found for the Ee- items (Ee- - Nee- = +3%). The lack of facilitation for Ee- items may have been due to convert retrieval on extra-exposure trials. If Ps were introducing suppression into these trials by engaging in covert retrieval, then predictions are less clear, for such retrievals could have led to suppression of Ee- items. Nonetheless, a simple linear regression analysis suggests that covert-retrieval activity, as measured by my questionnaire item, was not a significant predictor of Ee- facilitation (*beta* = -.015; R^2 = .004; F (1, 47) = .199, p = .658, B_{01} = .92). Anderson and Bell's (2001) method of assessing the influence of covert retrieval was also used. Specifically, the Ps were ranked according to covert retrieval scores within their respective counter-balance conditions. Ps ranking in the lower half of each condition were grouped together, as were Ps ranking in the upper half. No statistical analysis was required to see if covert-retrieval was diminishing the facilitation of Ee- items, for the high covert-retrieval Ps had greater numeric facilitation of Ee- items (3%) than did low covert-retrieval Ps (2%).

Experiment 2

The results of Experiment 1 failed to show support for a novel prediction of the *pattern-suppression model*, namely, the *pattern-facilitation hypothesis.* Experiments 2-3 tested predictions of the *pattern-suppression model* that have been confirmed in prior studies, but with a methodology that allowed us to see if such effects really are due to the mechanism of pattern suppression. That is, Experiment 2-3 allowed us to test if items suffering from *second-order inhibition* are only impaired with respect to features that overlap with Rp- items and *not* with respect to features that do not overlap with Rp- items.

In Experiment 2, Ps studied propositions (e.g., "The ant is crawling on the potato") as opposed to category-exemplar pairs. Each proposition (e.g., "The ant is crawling on the potato") can be decomposed into 2 parts: the topic ("The ant is crawling on the") and the object ("potato"). When comparing the propositions of Experiment 2 to the category-exemplar pairs of Experiment 1, the topic is analogous to the category and the object is analogous to the exemplar. No object was paired with multiple topics. RIF has been found when such propositions are used as stimuli (Anderson & Bell, 2001; Macleod & Saunders, 2006; Macrae & Macleod, 1999; Radvansky, 1999; Saunders & Macleod, 2005).

Unlike Experiment 1, there was only retrieval practice (and no extra exposure) in the second phase: half of the Ps did retrieval practice on half of the objects of *every* topic (*experimental* condition) and half of the Ps did retrieval practice only on fillers (*control* condition). Thus, there should only have been suppression of target objects in the *experimental* condition (see Figure 3). These objects were again termed Rp- items; the exact same objects in the *control* condition were termed Nrp items. Another change in Experiment 2 was that there was no final test that explicitly assessed memory for the studied material. Rather, Ps did a fact-verification task, in which Ps were presented with new propositions and had to determine as quickly as possible whether each proposition was a true fact of the world or was false. Half of the Ps were in the *overlap* condition, in which the predicates of each true fact (e.g., Chicken *is often baked*) was also true of an Rp-/Nrp item (e.g., potatoes are also *often baked*; see Figure 3). Each of these predicates thus acted as a semantic feature (e.g., *often baked*) embedded within the memory representations of both an Rp-/Nrp item (e.g., potato) and of a subject of a new proposition (e.g., chicken), and therefore was a point in which the two "semantically overlapped." The other half of the Ps were in the *no-overlap* condition, in which the predicates of the true facts (e.g., Chicken *is meat*) were *not* also true of the Rp-/Nrp items (see Figure 3).

The *pattern-suppression model* predicted that the overlapping features were themselves suppressed for Ps in the *experimental* condition. Consequently, the ability to activate these features within a memory representation should have been impaired, and so these Ps should have had increased difficulty in affirming these features (e.g., *is often baked*) as truly belonging to subjects (*chicken*) in the fact-verification task (see Figure 3). This predicted increase in difficulty was measured by comparing the performance (i.e., accuracy and speed) of Ps in the *experimental/overlap* group to the performance of Ps in the *control/overlap* group: the latter were predicted to perform better than the former on the fact-verification task. By contrast, Ps in the *experimental/no-overlap* group were not expected to perform worse than the *control/no-overlap* group on the fact-verification task. Hence, the *pattern-suppression model* predicted an interaction between the two variables.

Method

Participants

There were 216 undergraduates (67 women, 44 men, and 105 that did not report) from the University of Oklahoma who participated in Experiment 2. Ps were predominantly White or Caucasian (59%), with other Ps being American Indian or Alaskan Native (7%), Asian (7%), Black or African American (7%), Middle Eastern (1%), and Native Hawaiian or other Pacific Islander (1%); the remaining 10% chose "no response." Ps ranged in age from 18 to 33 years of age (M = 18.7). Ps received class credit for their participation.

Design

There were 2 independent variables (IVs), each of which was manipulated between-subjects. The first IV was the type of retrieval practice, which had two conditions: *experimental* and *control*. For Ps in the *experimental* condition, retrieval practice was over propositions of interest. In the *control* condition, there was only retrieval practice over fillers. The second IV was amount of overlap, which also had two conditions: *overlap* and *no-overlap*. In the *overlap* condition, Ps had to verify true facts that each had a predicate that was also true of an Rp-/Nrp item. In the *non-overlap* condition, Ps had to verify true facts that had predicates which were *not* also true of an originally studied object. RTs as well as accuracy were measured on the fact-verification task.

Procedure

Initial instruction and obtaining of verbal consent were the same as in Experiment 1. Ps were then instructed to study a series of 55 propositions (36 target propositions; 12 fillers propositions, 7 of which were shown twice). Ps were given 8 s to study each proposition. The propositions were presented in a randomized block, with each block containing a pairing of each topic with one of its objects. Furthermore, in order to minimize primacy and recency effects, the first 3 and last 3 propositions presented were fillers; no 2 topics appeared in sequence more than once; successive objects were never from the same topic; there was an average of 7 intervening trials in between successive objects from a topic; Rp+ (non-Nrp) and Rp- (Nrp) items were distributed evenly throughout the study phase.

After the initial study phase, Ps engaged in retrieval practice. Each trial provided a topic and the first 2 letters of an object (e.g., “The ant is crawling on the wi__”); Ps had 16 s to recall the appropriate object and to type it. Ps in the *experimental* condition did retrieval practice on 3 of the objects for each of the 6 topics. Furthermore, there were 3 retrieval-practice trials for each of these objects, with 3-4 intervening trials in between the first and second trial of a given object (average: 3.67) and 6 intervening trials in between the second and third trial of a given object. Lastly, there were no two objects of the same topic that were practiced successively and there were no repeating sequences. Filler items were used in order to meet these criteria. Ps in the *control* condition simply did retrieval practice on filler items.

After the retrieval-practice phase, Ps did a distractor task for 5 min. Ps were then told that they would be presented with a series of new propositions. These propositions had not been shown earlier and Ps had to indicate for each whether or not it was true of the real world, and to do so as quickly as possible. Ps indicated that a proposition was true by pressing the “j” button on the keyboard and that a proposition was false by pressing the “f” button.

Materials

Studied propositions consisted of two parts: a topic (e.g., “The ant is crawling on the “) and an object (e.g., “potato”). There were six topics, each of which was paired with 6 different objects for a total of 36 different propositions (see Appendix B). The studied propositions were largely the same as those used by Anderson and Bell (2001, Experiment 2). There were only two changes: the object *tulip* was replaced with *rose* and the object *vodka* was replaced with *rum*. These changes eased the process of producing propositions for the fact-verification task. No two objects belonged to the same category or to categories that are strongly related. The objects that received retrieval practice each began with a unique two-letter stem and none of the objects were associated with each other. Lastly, the propositions were so constructed as to not be memorable (i.e., distinctive or bizarre) or predictable.

All propositions in the fact-verification task (e.g., "Chicken is often baked") consisted of a subject (e.g., "Chicken") and a predicate (e.g., "is often baked), with none of the subjects associated with any of the objects of the originally studied propositions. The task began with 6 filler propositions (3 true, 3 false) in order to get Ps used to the routine. The true propositions in the *overlapping* condition were constructed such that each predicate was also true of one (and only one) Rp-/Nrp item. In order for the predicate of a true propositions to be considered as a semantic feature of an object, two conditions were met: the object (e.g., potato) had to be associated with the substantive term of the predicate (e.g., bake) and the predicate had to be true of the object (potatoes are indeed often baked). The same two conditions had to be met in order for the predicate of the Fact-Verification proposition to be considered as a semantic feature of the subject of that very same proposition (e.g., *Chicken*). Thus, when these two conditions were met in both instances, then the Rp-/Nrp item (*potato*) and the subject of the Fact-Verification proposition (*Chicken*) were considered to share the predicate as a semantic feature (*often baked*). In order to determine the existence of an association, I made recourse to the University of Southern Florida (USF) norms (Nelson, McEvoy, &

Schreiber, 1998). Each predicate in the *experimental* condition was a semantic feature of only one Rp-/Nrp item and of only one fact-verification subject. The true propositions for the *no-overlap* condition were constructed such that each predicate was neither associated with nor verily predicable of any of the objects from the originally studied set. However, the subjects (e.g., *Chicken*) of the propositions in the *no-overlap* condition were the same as those used in the *overlap* condition.

In order to prevent bias, half of the propositions of fact-verification task were false. The false propositions were constructed in the following way. First, true propositions were constructed along the same lines as were the true propositions in the *no-overlap* condition. Then the predicates were shuffled, resulting in new pairings, each of which was false. To reduce noise, these false propositions were constructed so as to be obviously false, and the true propositions were constructed as to be obviously true. Each participant responded to 18 true propositions and 18 false propositions. The propositions were presented in 6 blocks of 6 (3 true, 3 false).

The distractor task was the same as that used in Experiment 1.

Results

Retrieval-practice accuracy for Ps who did retrieval-practice on targets was .54, which was low when compared to the accuracies of other studies (see Table 2). This issue is addressed below. Test accuracy for the true facts in the fact-verification task was .93, suggesting that the truth values of these propositions were obvious, as was intended. However, Ps did have very poor accuracy (M = .55) for one fact from the *no-overlap* condition: *Ladders have rungs*. Hence, this proposition and its pair from the *overlap* condition, *Ladders let you move up*, were dropped from all subsequent analyses.

For many of the following analyses, the dependent variable was a combination of P's RT and accuracy in responding to the true facts in the fact-verification task. This combination dependent variable was termed *rate of correct responding* (RCR) and has precedent in the literature (e.g., Woltz, Sorensen, Indahl, & Splinter, 2015). The RCR score for each participant was calculated by dividing overall accuracy by RT (in min). A higher RCR score denotes a trend of responding more accurately and quickly. According to the *pattern-suppression model*, an interaction should be obtained in which the RCR scores of Ps in the *no-overlap condition* is not influenced by whether they are in the *experimental* condition or in the *control* condition; by contrast, the RCR scores of Ps in the *overlap* condition should be higher when those Ps are also in the *control* condition than when they are also in the *experimental* condition. A 2 (Overlap: Yes, No) × 2 (Retrieval Practice: Experimental, Control) analysis of variance (ANOVA) was run to test this prediction. First, there was no main effect of Overlap, F (1, 212) = 1.08, p = .18. However, there was a significant main effect of Retrieval Practice, F (1, 212) = 6.29, p = .01, but it appears to be due to the interaction, which was also significant, F (1, 212) = 4.99, p = .03 (see Figure 4a). Importantly, the nature of the interaction was counter to the prediction of the *pattern-suppression assumption*, with Ps in the *overlap* condition having greater RCR scores when in the *experimental* condition (M = 41.7, SD = 10.8) than when in the *control* condition (M = 35.6, SD = 35.6; see Figure 4). Hence, priming rather than suppression was observed. A post-hoc t-test revealed that the priming was significant, t (109) = 3.12**,** p = .002, Cohen's d = .60, B_{01} = .07. By contrast, for Ps in the *no overlap* condition, performance in the *experimental* condition (M = 37.1, SD = 9.0) and performance in the *control* condition (M = 36.7, SD = 7.8) did not differ, t (103) = .21, p = .83, Cohen's d = .05, B_{01} = 4.8.

The priming effect was surprising, and so additional analyses were performed to further explore the data. One important discovery was made when the foils (i.e., the false propositions) in the fact-verification task were made the dependent variable. In this analysis, the pattern of results was very similar to when the target items were the dependent variable (see Figure 4b). Importantly, however, the interaction was not significant, F (1, 212) = 2.5, p = .12. (The main effect of overlap was also not significant, F (1, 212) = .001, p = .98; however, the main effect of condition was significant, F (1, 212) = 5.3, $p < .05$ (= .02)). Nonetheless, Ps in the *overlap/experimental* group still seemed to perform better on the fact-verification task (M = 36.0, SD = 10.3) than did Ps in the *overlap/control* group (M = 31.4, SD = 8.9) [t (109) = 2.5, $p < .05$, Cohen's d = .48, B_{01} = .31], thereby suggesting that the initial priming that was observed may have been due merely to pre-existing differences in the groups rather than to manipulations that took place in the experiment.

A multiple regression analysis was therefore run to determine if RCR scores for true propositions in the *overlap* condition were still significantly predicted by whether one was in the *experimental* condition or in the *control* condition when RCR scores for foil items were taken into account. Hence, a hierarchical regression was run, in which foil RCR scores were inserted in the first step, followed by Retrieval-Practice condition in the second step. The first level of the analysis explained .728 of the variance, which was significant, F (1, 109) = 292.0, $p < .01$. However, the subsequent inclusion of retrieval practice condition as a predictor (beta = 1.90) still explained a marginally significant additional amount of the variance (change in R^2 = .008), F (1, 108) = 3.08, *p* = .082. Importantly, the analysis still suggests that retrieval-practice of Rp+ items results in, if anything, priming of Rp- features rather than suppression, as it still (numerically at least) boosted RCR scores by an average of 1.90.

The results suggest that performance for the Ps in *overlap*/*experimental* group was facilitated, whereas, the pattern-suppression model predicts impaired performance. There are at least three possible explanations for this discrepancy. One is that the *pattern-suppression model* is incorrect. This possibility will be explored in the Discussion section. Another possibility is that integration effects are responsible for the facilitation. Perhaps, during retrieval practice, Ps would recall the Rp- items covertly even though they were only instructed to recall the Rp+ items. Such activity has been found to counteract the effects of inhibition (Anderson & Bell, 2001). Unfortunately, data on integration were not taken, and so this possibility could not be explored very directly. However, RT during retrieval-practice could give a rough measure of integration activity, for such activity takes up time. Thus, if integration was responsible for the observed facilitation, then it would be predicted that Ps in the *overlap/experimental* group who spent more time on retrieval-practice trials would perform better in the fact-verification task. In analyzing the relationship between retrieval-practice RT and fact-verification performance, only correct responses from the retrieval-practice phase were taken into account, for what I was interested in was, when a participant thought he or she knew the answer on a trial, did he or she spend extra time thinking of other items. One issue was that greater RT's could also be indicative of poorer ability, in which case we would predict those with greater RT's to also perform worse (as opposed to better) on the fact-verification task. In other words, we would then hypothesize that greater RT's are predictive of both better performance on the fact-verification task (because of integration activity) as well as worse performance (because of poor ability). To reduce the noise provided by poor ability, I ran a hierarchical regression, on Ps in the *experimental/control* group, with two steps. In the first step, I analyzed how well retrieval-practice accuracy predicted performance on the fact-verification task. Retrieval-practice accuracy was meant to serve as an indicator of ability, thereby allowing us to partial out the effect of ability in the second step of the hierarchical analysis. The first level of the analysis explained .10 of the variance, which was significant, $F (1, 54) = 6.2$, $p < .05$. Correct retrieval-practice RT was inserted as a predictor (beta = -.005) in the second step of the analysis, and explained a significant additional amount of the variance (change in $R^2 = .20$), $F (1, 53) = 8.9$, $p < .005$. Importantly, the analysis suggests that longer correct retrieval-practice RT's result in worse (as opposed to better) performance on the fact-verification task for Ps in the *overlap/retrieval-practice* group, and so is inconsistent with the integration explanation. However, it should be remembered that RT is not an optimal measure for integration activity because it can be affected by other factors (e.g., additional rote rehearsal of Rp+ items).

The third possible explanation for the discrepancy between the data and the prediction of the *pattern-suppression model* involves the notion of suppression failure. Anderson and Levy (2011), based on a review of several interesting findings, suggest that a.) suppression attempts can fail and b.) if an attempt to suppress an item fails, then the accessibility of that item will be boosted. Knowing that suppression of competitor items is supposed to function as a means for facilitating memory, it could be argued that the low retrieval-practice accuracy (M = .54) is indicative of a significant amount of suppression failure. Furthermore, such suppression failure would result in facilitation of the Rp- items and, according to the *pattern-suppression model*, other items that have semantic overlap with the Rp- items, such as the subjects of the overlapping propositions in the fact-verification task. Consequently, if we assume that there was a significant amount of suppression failure, then the *pattern-suppression model* is then consistent with the finding that Ps in the *overlap* condition have higher (rather than lower) RCR scores if they are also in the *experimental* condition than if they are also in the *control* condition. In such a scenario, one prediction would be that, for *overlap/experimental* Ps, there should be a negative correlation between retrieval-practice accuracy and RCR scores in the fact-verification task, for greater accuracy is indicative of less suppression failure. However, the correlation was in fact positive, r = .32. This correlation could have been positive simply because low motivation/capability Ps scored low both in the retrieval-practice phase and in the fact-verification phase, and so a positive correlation was held within only their subgroup, but strongly affected the correlation of the whole group. Thus, the correlation was tracked as a function of retrieval-practice accuracy. As can be seen in Figure 5, once only Ps with the top 20 retrieval-practice accuracy scores (> .59 accuracy) were taken into account, the correlation had the predicted negative sign. A regression analysis with these 20 Ps did not, however, find retrieval-practice accuracy to be a significant predictor of performance on the fact-verification task (beta = -13.4), R_2 = .017, F (1, 18)

= .312, p = .58, but the subsequent Bayes Factor calculation did not provide evidence for the null hypothesis, B_{01} = .88. Hence, a final analysis was ran which compared the performance of these 20 Ps on the fact-verification task—who should have the least amount of suppression failure—to that of Ps in the *overlap/control* group. In order to give the most sensitive test for the suppression-failure explanation of the results, the analysis was restricted to the 20 Ps in the *overlap/control* group who scored the best on the fact-verification task; nevertheless, these Ps did not perform numerically as well on the fact-verification task (M = 46.4, SD = 6.4) as did the 20 Ps in the *overlap/retrieval-practice* group whom we have been considering (M = 47.1, SD = 11.5), although the difference was not significant, t (38) = .23, p = .82, B_{01} = 3.2. That is, even when we restrict the comparison to Ps in the *overlap/experimental* with the least amount of suppression failure and Ps in the *overlap/control* group with the best performances on the fact-verification task, we still do not find the performance of the former to be impaired relative to the performance of the latter. The suppression-failure explanation therefore does not seem to explain why suppression was not observed in the main analysis.

Experiment 3

I failed to find evidence for pattern-suppression as a mechanism for second-order inhibition in Experiment 2. However, as Radvansky (1999) pointed out, it may be useful to use a negative priming method. Such a method allows one to detect for suppression immediately after the suppression activity has occurred, which has two strengths. First, any suppression should not have yet dissipated naturally; second, there should be no intervening cognitive activity that disables the suppression (e.g., “release from inhibition," Bjork, Bjork, & Glenberg, 1973, as cited in Bjork, 1989) prior to its being measured on a fact-verification trial. For Experiment 3 of the current study, the negative priming method amounted to collapsing the retrieval-practice phase and the fact-verification phase together, such that after each retrieval-practice trial there was a fact-verification trial (see Figure 6). Notice that, in contrast to Experiment 2, the *control* condition does not consist of retrieval practice on fillers, but of retrieval practice on the same propositions receiving retrieval practice in the *experimental* condition. However, unlike in the *experimental* trials, the objects that are supposed to be suppressed do not share a major semantic feature with the subject of the ensuing fact-verification trial (see Figure 6). The prediction of the *pattern-suppression model* was the same for Experiment 3 as it was for Experiment 2: fact-verification performances in the *experimental/overlap* condition should be worse than performances in the *control/overlap* condition, but the performances of the *experimental/no-overlap* condition and of the *control/no-overlap* condition should not differ.

Method

Participants

There were 84 undergraduates (68 women) from the University of Oklahoma who participated in Experiment 3. The majority of the Ps were White or Caucasian (75%), with the other Ps being Black or African American (5%), American Indian or Native Alaskan (4%), Middle Eastern (1%), Native Hawaiian or other Pacific Islander (1%), or giving "no response" (12%). Ps ranged in age from 18-22 (M = 18.7). Ps received class credit for their participation.

Design

There were two within-subjects IVs, each with 2 conditions. The first IV was amount of overlap, with the two conditions being *overlap* and *no-overlap*. This IV concerns the true propositions in the fact-verification trials. The true propositions had the same subject in both conditions (e.g., Chicken), but the predicate for each subject differed between the conditions (e.g., “Chicken is often baked” for the *overlap* condition and “Chicken is meat” for the *no-overlap* condition). In the *overlap* condition, the predicate was also true of an object from one of the originally studied propositions (e.g., potatoes are also *often baked* and Ps originally studied “The ant is crawling on the potato.”); in the *no-overlap* condition, the predicate was not also true of an object from one of the originally studied propositions (e.g., potatoes are *not* also meat).

The second IV was prime type, which concerns the retrieval-practice trials, and its two conditions were termed *experimental* and *control*. As in Experiment 2, the true propositions in the fact-verification task each had a subject that had semantic overlap with one of the objects from the initially studied propositions. In the *experimental* condition, the subject of a fact-verification task proposition had semantic overlap with an object that should have been suppressed during the immediately preceding retrieval-practice trial. In the *control* condition, the subject of a fact-verification task proposition did not have semantic overlap with any of the objects that were supposed to have been suppressed in the immediately preceding retrieval-practice trial (see Figure 6).

Procedure

Initial instruction and obtaining of verbal consent were the same as the previous experiments; the initial study phase was no different than that of Experiment 2. Ps were then given instructions for the final test, which included retrieval-practice trials as well as fact-verification trials, with each retrieval-practice trial being followed immediately by a fact-verification trial. Ps were next questioned about their study strategies before being thanked for their participation.

Materials

Both the studied propositions and the fact-verification propositions were the same as those used in Experiment 2. However, additional fillers study propositions were added, as were additional fact-verification propositions, half of which were true and half of which were false. These latter propositions were constructed in the same way as in Experiment 2, and were used in the final trial of each *micro-block* (which are explained below).

On the final test, every retrieval-practice trial was followed by a fact-verification trial, and, for simplicity, each such set of 2 trials can be termed a *unit*. (The following design is quite complex. See Figure 7 for an illustration). Each retrieval-practice trial presented a topic along with the first 2 letters of one its objects (e.g., "The ant is crawling on the po__"). In order to allow Ps to practice, the first 6 units of the final test were fillers. After those initial 6 units, Ps completed the remaining units in 4 blocks. However, each of these blocks consisted of 3 smaller blocks; the former were termed *macro-blocks* and the latter *micro-blocks.* Each *micro-block* consisted of 7 units, the last of which was always a filler unit (in this case, retrieval-practice and fact-verification were both over filler propositions, half of which were true). The remaining 6 units each had a retrieval-practice trial that corresponded to a unique topic (e.g., "The ant was crawling on the"). The filler unit was always presented last in each *micro-block* to ensure that, in the transition from one *micro-block* to another, there were not successive retrieval-practice trials corresponding to the same topic. The order of the remaining units within a *micro-block* was random. Ps did retrieval practice on 3 of the objects for each of the topics; each such object received retrieval-practice in one and only one *micro-block* of each *macro-block*, and hence four times total. Whether a given object received retrieval practice in the first, second, or third *micro-block* was determined randomly within each of the four *macro blocks.* In two of the *macro blocks*, retrieval-practice of a given object was followed by a fact-verification trial over a false proposition (foil). These false propositions were used in order to prevent a bias to respond "true" on the fact-verification trials. The retrieval-practice of a given object was followed by a fact-verification foil once within the first two *macro blocks* and once within the last two *macro blocks*; further specification of the two *macro blocks* within which retrieval-practice of an object would be followed by a fact-verification foil was random. In the other two *macro blocks*, retrieval-practice of that object was followed by a true fact-verification trial once in the *experimental* condition and once in the *control* condition (see Table 3, Figure 7). For all Ps, half of the objects receiving retrieval practice were followed by an *experimental/overlap* fact-verification trial and by a *control*/*overlap* fact-verification trial, whereas the other half of the objects receiving retrieval practice were followed by an *experimental/no-overlap* fact-verification trial and by a *control/no-overlap* fact-verification trial (see Table 3, Figure 7).

In order to rule out confounding factors, 4 counter-balancing groups were constructed (see Table 3). In the following, all reference to topic order (e.g., what the "first three" topics are) are in accordance with a set order. Hence, the "first" topic was the same for all Ps, although it may not have been the first topic encountered by a given. In the first and third counter-balancing group, the first 3 topics made up the *units* in the *overlap* condition and the *last* 3 topics made up the *units* in the *no-overlap* condition; the opposite was the case for the second and fourth counter-balancing group. In the first and second counter-balancing group, *experimental* trials were done before the *control* trials for the first 3 topics, but after for the last 3 topics; the opposite was the case for the third and fourth counter-balancing group.

Results

Retrieval-practice accuracy was .52, which was low when compared to the accuracies of other studies (see Table 2). This issue is addressed below. Test accuracy for the true facts in the fact-verification task was .94, suggesting that the truth values of these propositions were obvious, as was intended. However, Ps did again have very poor accuracy (M = .63) for the fact *Ladders have rungs*. Hence, this proposition and its pair from the overlap condition, *Ladders let you move up*, were dropped from all subsequent analyses.

The *pattern-suppression model* predicts an interaction in which, for the *overlap condition,* performance on *control* fact-verification trials should be worse than performance on *experimental* trials; by contrast, performance on *control* trials should not differ from performance on *experimental* trials in the *no-overlap* condition. The data were analyzed using a 2 (Overlap: Yes, No) × 2 (Prime Type: Control, Experimental) ANOVA. There was neither a main effect of Overlap, F (1, 83) = 1.87, *p* = .18, nor of Retrieval-Practice, F (1, 83) = .44, *p* = .51; there was only a marginally significant interaction, *F* (1, 83) = 2.80, *p* = .10. It should be noted, however, that in the *overlap condition*, in contrast to the prediction of the *pattern-suppression model*, there were numerically higher RCR scores on *experimental* trials (*M* = 32.7, *SD* = 9.2) than on *control* trials (*M* = 32.3, *SD* = 10.1; see Figure 8), although analyses strongly suggest that the difference is illusory, *t* (83) = -.581, *p* = .56, Cohen's *d* = .04, B_{01} = 7.1. In the *no overlap* condition, there was also no significant difference between performance in the *control condition* (*M* = 31.9, *SD* = 10.1) and performance in the *retrieval-practice condition* (*M* = 31.0, *SD* = 10.1), *t* (83) = 1.5, *p* = .14, Cohen's *d* = .09, B_{01} = 2.8.

Accuracy on the retrieval-practice trials was .52, which again is low and suggests the possibility that suppression failure is the cause of the numerical facilitation (as opposed to the predicted suppression) in the fact-verification. As argued in the results section of Experiment 2, if suppression failure is the reason why the original prediction of the *pattern-suppression model* is violated, then, for the *experimental/overlap* units, there should be a negative correlation between retrieval-practice accuracy and fact-verification performance. However, such a negative correlation would be blurred by the performance of low ability/motivation Ps: these Ps are likely to perform very poorly on both the retrieval-practice trial and on the fact-verification trial. To reduce this noise, the dependent variable was changed from performance on *experimental/overlap* fact-verification trials to a difference score:

[*experimental/overlap* trial performance] – [*control/overlap* trial performance]. The predictor variable was retrieval-practice accuracy on *experimental/overlap* retrieval-practice trials. As predicted, the overall observed relationship was negative (*unstandardized* beta =-2.4; *standardized* beta = -.11), but it was not significant F (1, 82) = .98, p = .325 (R^2 = .012). Hence, I failed to find evidence that the effects of suppression were being masked by suppression failure.

Another possibility is that integration activity was occurring. As in Experiment 2, I again used RT on correct retrieval-practice trials as a proxy for integration activity, with greater RT (on accurate trials) marking greater integration activity. If this is the explanation for why suppression was not detected, then greater *experimental/overlap* retrieval-practice RT (i.e., greater integration activity) should be positively correlated with the following difference score: [*experimental/overlap* trial performance] – [*control/overlap* trial performance]. However, as in Experiment 2, I again wanted to partial out the effect of motivation/ability on *experimental/overlap* retrieval-practice RT, and so again ran a hierarchical regression analysis. In the first step, difference scores were regressed onto retrieval-practice accuracy, which served as indicator of motivation/capability. Retrieval-practice accuracy explained .002 of the variance, which was not significant, F (1, 81) = 174, p = .68. In the second step, I saw if adding *experimental/overlap* retrieval-practice RT for correct trials significantly improved the regression, but it did not [change in R^2= .03, F (1, 80) = 2.3, p = .14 (*unstandardized* beta = .001; *standardized* beta = .166)]. Hence, I did not find evidence that integration activity was the reason why suppression was not detected, although the analysis should be interpreted with caution because retrieval-practice RT is only a rough indicator of integration activity.

Discussion

In the current study, I have run 3 experiments, each of which tested a prediction of the *pattern-suppression model*. In Experiment 1, it was predicted that items in memory would become more accessible if Ps studied other items that had high semantic overlap with them. However, no such increased accessibility was detected. In Experiment 2, I looked at the accessibility of specific features of items in memory. According to the *pattern-suppression model*, Ps should have had impaired access to a specific subset of those features; however, Ps had facilitated access to that subset of features. In Experiment 3, Ps had numerically facilitated access to a specific subset of features of the items that, according to the *pattern-suppression model*, should have been suppressed. I will give a theoretical analysis for each of these findings in turn.

In Experiment 1, extra-exposure of items did not lead to increased or decreased accessibility of other items that had high semantic overlap with them. This finding is consistent with past studies on RIF (e.g., Anderson & Bell, 2001; Anderson et al., 2000; Bäuml, 2002), although there was not very high semantic overlap between Ee+ items and Ee- items in those studies. The reason why there was not high semantic overlap between Ee+ items and Ee- items in those studies was that the subject of concern in them was the *retrieval-specificity assumption* of the suppression theory and not the *patter-suppression model*. According to the *retrieval specificity assumption*, additional encoding of a memory item only results in impairment of related items if that additional encoding is the result of memory retrieval and *not* if it is a non-retrieval form of encoding (e.g., studying a *presented* word). This assumption is based on the idea that, because suppression is a mechanism used to aid retrieval and is of no value to merely processing that which is already present, suppression will not operate in the processing of that which is already present. Hence, the lack of facilitation in those past studies was not viewed as evidence against the suppression theory; rather, resulting lack of impairment was viewed as a prediction of the suppression theory. Of course, if there was *any* overlap between the Ee+ items and the Ee- items, there should have been *some* facilitation of the Ee- items. As explained earlier, however, if the overlap was relatively little, then the facilitation may have been there but not detectable. The point of having high overlap between the Ee+ items and the Ee- items in Experiment 1 was to make such facilitation more detectable; nonetheless, (if present at all) it still went undetected.

Research from the *semantic-priming* literature, however, might shed light on the results from Experiment 1. *Semantic priming* occurs when the processing of one item facilitates the processing of a semantically related item. Dannenbring and Briand (1982), for example, found evidence of semantic priming in Ps who were executing a lexical-decision task. In each trial of the task, Ps would be shown a string of letters and would have to state if the string was a word or not. It was found that participants were more quickly able to determine that a string of letters was a word (e.g., ocean) if a previous trial was over a semantically related word (e.g., sea)—this finding constituted evidence of semantic priming. However, semantic priming *only* occurred when the two semantically related words were presented on consecutive trials; if trials not containing a semantically related word were placed between the two, then the priming effect dissipated. Thus, semantic priming may have occurred in Experiment 1, but dissipated by the time of the final test.

However, the results of Dannenbring and Briand (1982) alone would leave unexplained why, in both Experiment 1 and in Goodmon and Anderson (2011), retrieval-practice of Rp+ items resulted in no suppression of semantically overlapping Rp- items. According to the *pattern-suppression model*, the features of Rp- items that do not overlap with Rp+ items should have been suppressed, thereby making retrieval of the Rp- items during the final test more difficult. The reason why no impairment was found, according to the *pattern-suppression model*, was that the features of the Rp- items that did overlap with the Rp+ items were facilitated (or primed) and thus offset the impairment. However, such priming has often been found to be very temporary (Dannenbring & Briand, 1982) and so should have dissipated by the time of the final test, meaning that impairment should have been observed.

Recent work on semantic priming has found that it can be durable *if* the processing of the prime is very semantic in nature (Hughes & Whittlesea, 2003; Woltz et al., 2015). Determining whether it is a peel or a wing that is a feature that belongs to an *apple* requires greater processing into the *meaning* of *apple* than does determining that a string of letters spells the word *apple*. Likewise, using a category (e.g., *FRUIT*) as a cue to retrieve a particular exemplar (e.g., *apple*) of that category may require greater semantic processing than merely reading a category-exemplar pair. The former *requires* significant processing of the categorical features of the exemplar; the latter may consist more of phonological processing (i.e., how do the words sound) or other types of processing, resulting in facilitation of non-semantic features and, hence, features not necessarily shared with the other, semantically-overlapping exemplars. Hence, the results of Experiment 1 can, according to the *pattern-suppression model*, be explained in this way: retrieval-practice of an Rp+ item results in the suppression of non-overlapping features in Rp- items and in the durable semantic priming of overlapping features in Rp- items. The net result was therefore neither facilitation nor impairment on the final test. Extra exposure of Ee+ items results in neither suppression of non-overlapping features of Ee- items nor in durable priming of the overlapping features of Rp- items. The result again is therefore neither facilitation nor impairment on the final test. In order to test this explanation of Experiment 1, a future study should be done in which extra-exposure is replaced with a feature-selection task in which the exemplar is presented, but is flanked by two words, one of which is a feature of the exemplar and one of which is not (e.g., *peel-apple-wing*). The P's task would be to select the feature (*peel*) that belongs to the exemplar (*apple*). Such a task would require greater semantic processing, and has been found to result in durable semantic priming (Hughes & Whittlesea, 2003). Furthermore, there would be no need in such a task to suppress the other exemplars belonging to the same category. Hence, the durable priming of and lack of suppression of other exemplars from the same category should, according to the *pattern-suppression model*, result in increased accessibility of those other exemplars.

The results of Experiment 2 and 3 are more difficult to align with the *pattern-suppression model*. The procedure used to induce suppression of Rp- items (i.e., retrieval-practice) and the stimuli that were studied already have precedent in the RIF literature (Anderson & Bell, 2001, Experiment 2). Features (e.g., *in the woods*) of the

Rp- items (e.g., *tent*) that did not overlap with Rp+ items should have therefore been suppressed. Nonetheless, Ps were unimpaired in their ability to correctly affirm these very features as belonging to the novel items presented in the Fact-Verification task (e.g., *cabins*); in fact, in the main analysis of Experiment 2, P's ability to do so appeared to be facilitated. Ancillary analyses found no support for alternative, suppression explanations, such as suppression failure and integration, for the lack of observed suppression. Furthermore, although the integration explanation could only be tested in a rough manner, it should be kept in mind that spontaneous integration effects are not common in the RIF literature—most studies detect impairment just fine without taking integration activity into account. The greatest concern was that, in Experiment 2, group performance on foils—which should have not varied among the 4 groups—resembled group performance on target items. Nonetheless, an ancillary analysis found that, even when taking this issue into account, retrieval-practice still led to at least numerical facilitation rather than to the predicted impairment.

There are a few ways one can approach the data from Experiments 2 and 3. First, one can see them as a first attempt to test for suppression at the feature-level and, as such, merely suggestive rather than compelling. Second, one could look at the data as compelling evidence against the *pattern-suppression model*. After all, if anything, there was facilitation rather than suppression of features that, according to the model, should have been suppressed. If the *pattern-suppression model* is incorrect, then the finding of *second-order inhibition* is left currently inexplicable within the suppression framework (and within any other current framework) and is in need of further exploration. One important point to stress here is that *second-order inhibition* is a controversial finding that has been hard to replicate (Raaijmakers & Jakab, 2013; Perfect et al., 2014; Williams & Zacks, 2001) and that is not predicted by competition-based models of forgetting, such as the SAM-REM model (Malmberg & Shiffrin, 2005)

Lastly, one can look at the data as informative of the recently growing literature on durable semantic priming. One of the current debates is over the nature of semantic priming. Woltz et al. (2015) argue that the semantic *content* of an item is activated/excited during semantic processing of that item and that, therefore, *any* subsequent semantic processing of items that share that content will be facilitated. This theory of priming is very similar to Anderson's *pattern-suppression model*, or at least to its prediction that there is facilitation of features in Rp- items that overlap with Rp+ items. Furthermore, it is consistent with any facilitation that may have occurred in this study.

Alternatively, Hughes and Whittlesea (2003) have a very different account of semantic priming. They believe that semantic priming can only occur to the degree that a particular *operation* of semantic processing matches a previous operation of semantic processing. For example, determining that a *peel* is a feature of an *apple* will facilitate one's ability to affirm that a *peel* is a feature of an *orange*, but will not facilitate one's ability to affirm that *oranges* and *pears* belong to the same category. Likewise, determining that a *pear* and an *apple* are in the same category will facilitate one's ability to affirm that a *pear* and an *orange* are also in the same category, but it will not facilitate one's ability to affirm that a *peel* is a feature of an *orange*. For Hughes and Whittlesea (2003), the overlap must be in the semantic processing operations required rather than merely in the semantic contents of the objects that are semantically processed.

The overlapping-operations framework of Hughes and Whittlesea (2003) cannot account for any semantic priming that occurred in the current experiments. In Experiment 1, Ps may have had no detectable impaired ability to retrieve Rp- items on the final test because there was priming to counteract it. In the retrieval-practice phase, Ps had to process the exemplars of a given category in order to remember which of its exemplar began with a certain two-letter stem. During the final test, Ps have to do the same memory-search/processing of the same categories in order to remember the Rp- items. Hence, it appears at first that overlapping operations may be the cause of the hypothesized priming. However, these overlapping operations are also present in many (indeed most) RIF studies, in which the impairment of Rp- items *is* detectable and hence appears to *not* be counteracted by any semantic priming. The determinant of whether impairment is detected or not appears to be whether the semantic *contents* of the Rp- items overlap substantially with the Rp+ items. Only when there is substantial semantic overlap between Rp- items and Rp+ items does the impaired access to the Rp- items get counteracted, a finding more consistent with Woltz et al.'s (2015) theory of semantic priming than with Hughes and Whittlesea's (2003).

Any priming in Experiments 2 and 3 would also be inconsistent with Hughes and Whittlesea's (2003) overlapping-operations account. However, the first question is why there would be any priming in these experiments at all. The primes in these experiments would have been Rp- items, and the priming has to occur during the retrieval-practice trials. Ps were not asked to consciously recall Rp- (e.g., *tents*) items or to process them at all during the retrieval-practice phase; rather, Ps were recalling the Rp+ items. Yet, it appears that Ps may have been primed in their ability to affirm features of the Rp- items (e.g., *in the woods*) as belonging to novel stimuli (e.g., *cabin*). A possible explanation for this odd finding can be found in Carpenter's (2009) *elaborative-retrieval hypothesis*, which was developed to explain the *testing effect*. The *testing effect* refers to the fact that actively retrieving information from memory is a more potent form of encoding that information than is merely being exposed to that information. For example, one will be better able to *later* remember the definition of *defenestrate* if one *now* recalls it from memory than if one merely looks it up. According to the *elaborative-retrieval hypothesis*, when one is to use a cue (e.g., *defenestrate*) in order to remember a particular item (its definition), one will do a memory search through items that are associated to the cue in the memory network. Items that are activated during the search become associated both to the cue and to the sought item (if the sought item is either correctly recalled or provided via feedback). Hence, according to Carpenter (2009), the *testing effect* is a result of a more elaborate network of paths connecting the cue to the sought item. For our purposes, the theory suggests that when a *semantic* cue is being used, items (such as the Rp- items) may be activated via a semantic processing, thereby allowing for semantic priming of other items that overlap with them. In this manner, retrieval-practice in Experiments 2-3 may have resulted in priming of overlapping, true propositions in the fact-verification task.

If there was indeed priming in Experiment 2 and it is to be accounted for by the *elaborative retrieval hypothesis*, then two things follow. First, semantic priming in Experiments 2 would be inconsistent with the overlapping-operations account of semantic priming, for (as an example) determining whether or not *the keys are in the tent* (an Rp- proposition) is a very different process from determining whether or not *cabins are in the woods* (a fact-verification proposition). The second thing that follows is that the suppression theory of forgetting is incorrect, for it is in flat contradiction to the *elaborative-retrieval hypothesis*. The major items associated to a cue that are not the sought item must be suppressed in order to remember the sought item; hence, any connections from the cue to the sought item that involve these major items will be *impaired* rather than facilitated as a result of retrieval, according the suppression theory of forgetting.

References

Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language, 49*(4), 415-445. doi:http://dx.doi.org/10.1016/j.jml.2003.08.006

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(5), 1063-1087. doi:http://dx.doi.org/10.1037/0278-7393.20.5.1063

Anderson, M. C., Bjork, E. L., & Bjork, R. A. (2000). Retrieval-induced forgetting: Evidence for a recall specific mechanism. *Psychonomic Bulletin & Review, 7*(3), 522-530.

Anderson, M. C., & Levy, B. J. (2011). On the relationship between interference and inhibition in cognition. *Successful remembering and successful forgetting: A festschrift in honor of Robert A. Bjork.* (pp. 107-132) Psychology Press, New York, NY.

Anderson, M. C., Green, C., & McCulloch, K. C. (2000). Similarity and inhibition in long-term memory: Evidence for a two-factor theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(5), 1141-1159. doi:http://dx.doi.org/10.1037/0278-7393.26.5.114

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review, 102*(1), 68-100. doi:http://dx.doi.org/10.1037/0033-295X.102.1.68

Bäuml, K. (2002). Semantic generation can cause episodic forgetting. *Psychological Science, 13*(4), 356-360. doi:http://dx.doi.org/10.1111/j.0956-7976.2002.00464.x

Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. *Varieties of memory and consciousness: Essays in honour of endel tulving.* (pp. 309-330) Lawrence Erlbaum Associates, Inc, Hillsdale, NJ.

Broadbent, D. E., Cooper, P. F., FitzGerald, P., & Parkes, K. R. (1982). The cognitive failures questionnaire (CFQ) and its correlates. *British Journal of Clinical Psychology, 21*(1), 1-16.

Brown, R., & McNeill, D. (1966). The "tip of the tongue" phenomenon. *Journal of Verbal Learning & Verbal Behavior, 5*(4), 325-337.

Carpenter, S. K. (2009). Cue strength as a moderator of the testing effect: The benefits of elaborative retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*(6), 1563-1569. doi:http://dx.doi.org/10.1037/a0017021

Dannenbring, G. L., & Briand, K. (1982). Semantic priming and the word repetition effect in a lexical decision task. *Canadian Journal of Psychology/Revue Canadienne De Psychologie, 36*(3), 435-444. doi:http://dx.doi.org/10.1037/h0080650

Goodmon, L. B., & Anderson, M. C. (2011). Semantic integration as a boundary condition on inhibitory processes in episodic retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 37*(2), 416-436. doi:http://dx.doi.org/10.1037/a0021963

Hughes, A. D., & Whittlesea, B. W. A. (2003). Long-term semantic transfer: An overlapping-operations account. *Memory & Cognition, 31*(3), 401-411. doi:http://dx.doi.org/10.3758/BF03194398

Jonker, T. R., Seli, P., & MacLeod, C. M. (2013). Putting retrieval-induced forgetting in context: An inhibition-free, context-based account. *Psychological Review, 120*(4), 852-872. doi:http://dx.doi.org/10.1037/a0034246

MacLeod, M. D., & Saunders, J. (2005). The role of inhibitory control in the production of misinformation effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(5), 964-979. doi:http://dx.doi.org/10.1037/0278-7393.31.5.964

MacLeod, M. D., & Macrae, C. N. (2001). Gone but not forgotten: The transient nature of retrieval-induced forgetting. *Psychological Science, 12*(2), 148-152. doi:http://dx.doi.org/10.1111/1467-9280.00325

Macrae, C. N., & MacLeod, M. D. (1999). On recollections lost: When practice makes imperfect. *Journal of Personality and Social Psychology, 77*(3), 463-473.

Malmberg, K. J., & Shiffrin, R. M. (2005). The "one-shot" hypothesis for context storage. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(2), 322-336.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms. http://www.usf.edu/FreeAssociation/.

Perfect, T. J., Stark, L., Tree, J. J., Moulin, C. J. A., Ahmed, L., & Hutter, R. (2004). Transfer appropriate forgetting: The cue-dependent nature of retrieval-induced forgetting. *Journal of Memory and Language, 51*(3), 399-417.

Raaijmakers, J. G. W., & Jakab, E. (2013). Rethinking inhibition theory: On the problematic status of the inhibition theory for forgetting. *Journal of Memory and Language, 68*(2), 98-122. doi:http://dx.doi.org/10.1016/j.jml.2012.10.002

Radvansky, G. A. (1999). Memory retrieval and suppression: The inhibition of situation models. *Journal of Experimental Psychology: General*, *128*(4), 563-579. doi:http://dx.doi.org/10.1037/0096-3445.128.4.563

Saunders, J., & MacLeod, M. D. (2006). Can inhibition resolve retrieval competition through the control of spreading activation? *Memory & Cognition*, *34*(2), 307-322. doi:http://dx.doi.org/10.3758/BF03193409

Storm, B. C., & Levy, B. J. (2012). A progress report on the inhibitory account of retrieval-induced forgetting. *Memory & Cognition*, *40*(6), 827-843. doi:http://dx.doi.org/10.3758/s13421-0120211-7

Storm, B. C., & Nestojko, J. F. (2010). Successful inhibition, unsuccessful retrieval: Manipulating time and success during retrieval practice. *Memory*, *18*(2), 99-114.

Terman, M, Rifkin, J. B., Jacobs, J., & White, T. M. (2001). New York State Psychiatric Institute, 1051 Riverside Drive, Unit 50, New York, NY, 10032.

Verde, M. F., & Perfect, T. J. (2011). Retrieval-induced forgetting in recognition is absent under time pressure. *Psychonomic Bulletin & Review*, *18*(6), 1166-1171. doi:http://dx.doi.org/10.3758/s13423-011-0143-4

Williams, C. C., & Zacks, R. T. (2001). Is retrieval-induced forgetting an inhibitory process? *The American Journal of Psychology*, *114*(3), 329-354.

Woltz, D. J., Sorensen, L. J., Indahl, T. C., & Splinter, A. F. (2015). Long-term semantic priming of propositions representing general knowledge. *Journal of Memory and Language*, *79-80*, 30-52. doi:http://dx.doi.org/10.1016/j.jml.2014.11.002

Tables

Table 1

Experiment 1 Final Test Performance

Item Type	Percent Recall	Baseline Performance	Difference
Rp+	.74	.51	+.23*†
Rp-	.44	.52	-.08
Ee+	.66	.56	+.10*†
Ee-	$.59_a$	$.56_a$	+.03

Note: Means sharing a subscript were found to have a significant lack of difference at $B_{01} = 3$ according to Bayesian analysis.

* $p < .05$

† $B_{01} < 1/3$

Table 2

Retrieval-Practice Accuracy in Past Studies

Study	Retrieval-Practice Accuracy (*M*)
Anderson, Bjork, & Bjork (1994)	.68 - .90
Anderson & Spellman (1995)	.69 - .78
Anderson & McCulloch (1999)	.86 - .95
Anderson, Bjork, and Bjork (2000)	.83
Anderson, Green, and McCulloch (2000)	.73 - .81
Anderson & Bell (2001)	.62 - .80
Goodmon & Anderson (2011)	.87 - .94

Note: Most of these studies consisted of multiple experiments, and so the range of retrieval-practice accuracies are reported for each such study. Anderson and Bell (2001) was the only study in which Ps studied non-semantic propositions (like in my Experiment 2); the others worked with category-exemplar pairs, which are presumably more easy to remember.

Table 3

Experiment 3 Counter-Balancing Groups

Group	Topics	Macro Blocks	Prime Type	Overlap
1	1-3	1-2	Experimental	Yes
		3-4	Control	Yes
	4-6	1-2	Control	No
		3-4	Experimental	No
2	1-3	1-2	Experimental	No
		3-4	Control	No
	4-6	1-2	Control	Yes
		3-4	Experimental	Yes
3	1-3	1-2	Control	Yes
		3-4	Experimental	Yes
	4-6	1-2	Experimental	No
		3-4	Control	No
4	1-3	1-2	Control	No
		3-4	Experimental	No
	4-6	1-2	Experimental	Yes
		3-4	Control	Yes

Note: Each object that received retrieval-practice was practiced once in each *macro-block*. It is important to note that, in two of the *macro-block*, a given item was succeeded by a fact-verification trial over a false proposition (filler).

Figures

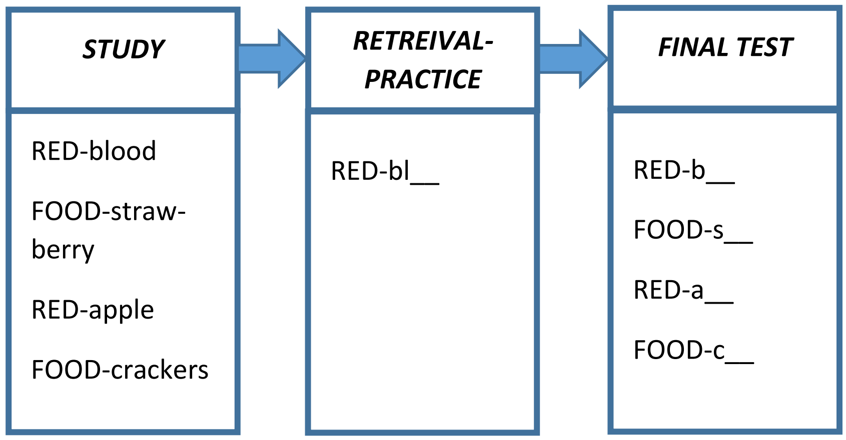

Figure 1. Representation of the retrieval-practice paradigm, as used by Anderson and Spellman (1995). Ps first study category-exemplar pairs. Next, they do retrieval practice for half of the exemplars (*blood*) of half of the categories (*RED*). These exemplars are termed Rp+ items; exemplars belonging to the same category as Rp+ items but that do not receive retrieval practice (*apple*) are Rp- items. Exemplars that do not receive retrieval practice and that do *not* belong to the same category as an Rp+ item are termed Nrp items (*strawberry; crackers*), which serve as a baseline against which to compare the success of recalling Rp+ and Rp- items during the final test. Often, the Nrp items will be divided into two groups, with one group—the Nrp+ items—serving as a baseline for Rp+ items and with the other group—the Nrp- items—serving as a baseline for the Rp- items.

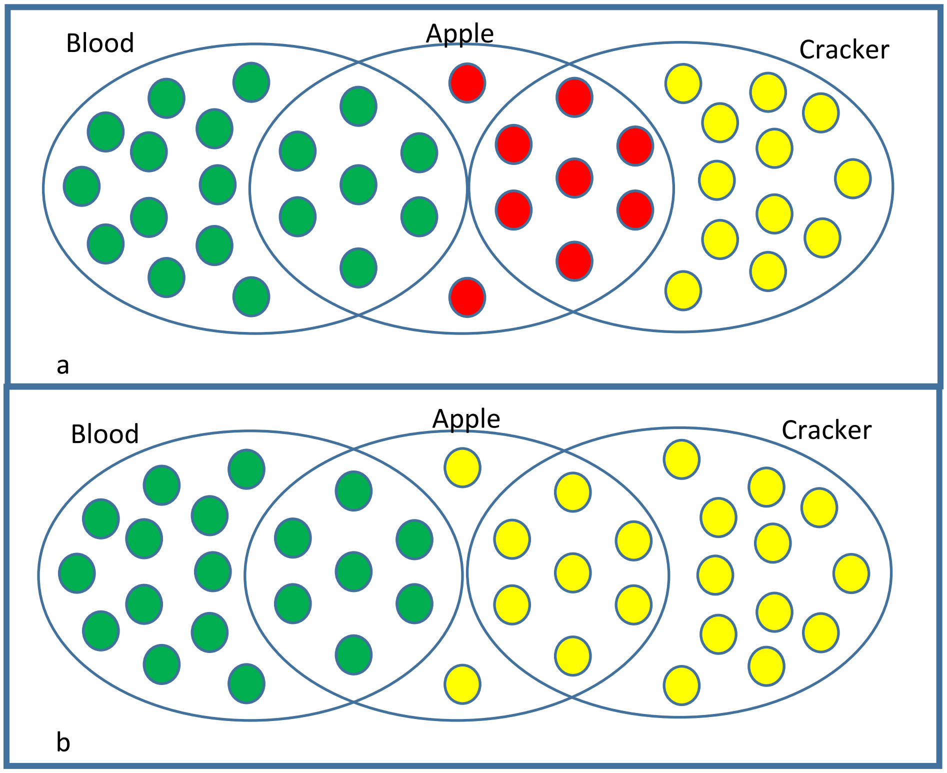

Figure 2. Pictorial representation of the *pattern-suppression model*. The model predicts that attempt to remember an item (*blood*) will have the following interesting effects that are displayed in (a). First, only the features (*fruit, round, etc.*) of the competitor (*apple*) that are not shared with the target (*blood*) will be suppressed (red circles). Furthermore, if the target item is recalled, then features that the target and competitor share (*red, material, etc.*) will become active in the memory representations of both (green circles). Hence, the degree to which the competitor is suppressed coincides with the degree to which it overlaps with the target: greater overlap will result in less suppression. Lastly, non-competitors (*crackers*) which are not directly suppressed can still be impaired to the extent that they share features with a competitor (*apple*) that have been suppressed (*food features*). We see in (b) what may follow from the exposure of a target as opposed to its retrieval. Because no suppression is required, overlapping items will no longer get any of their features suppressed. However, features they have that overlap with the target will still be facilitated.

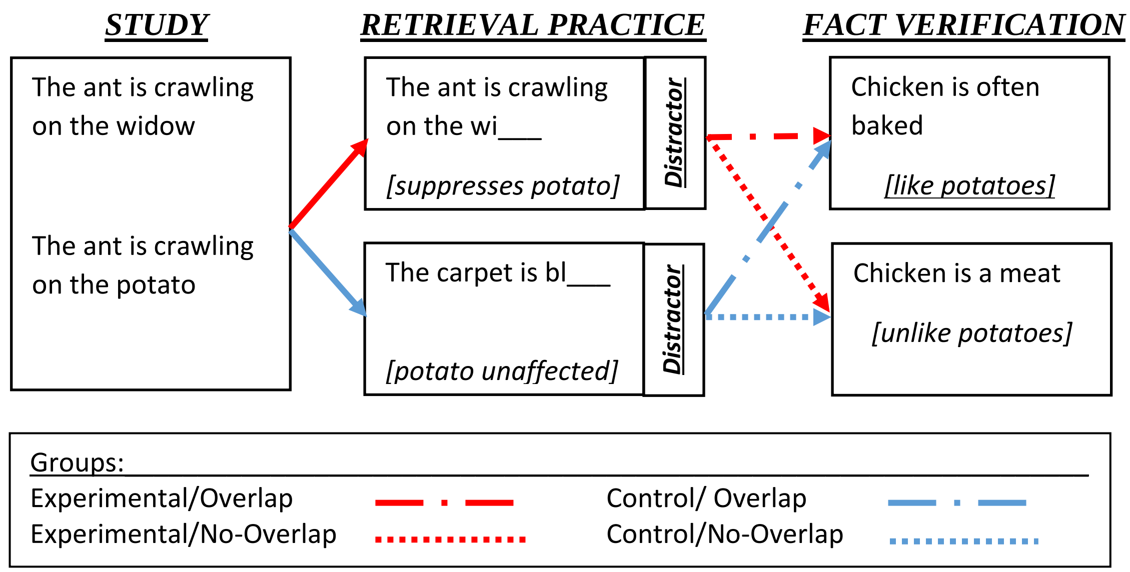

Figure 3. Design Scheme for Experiment 2.

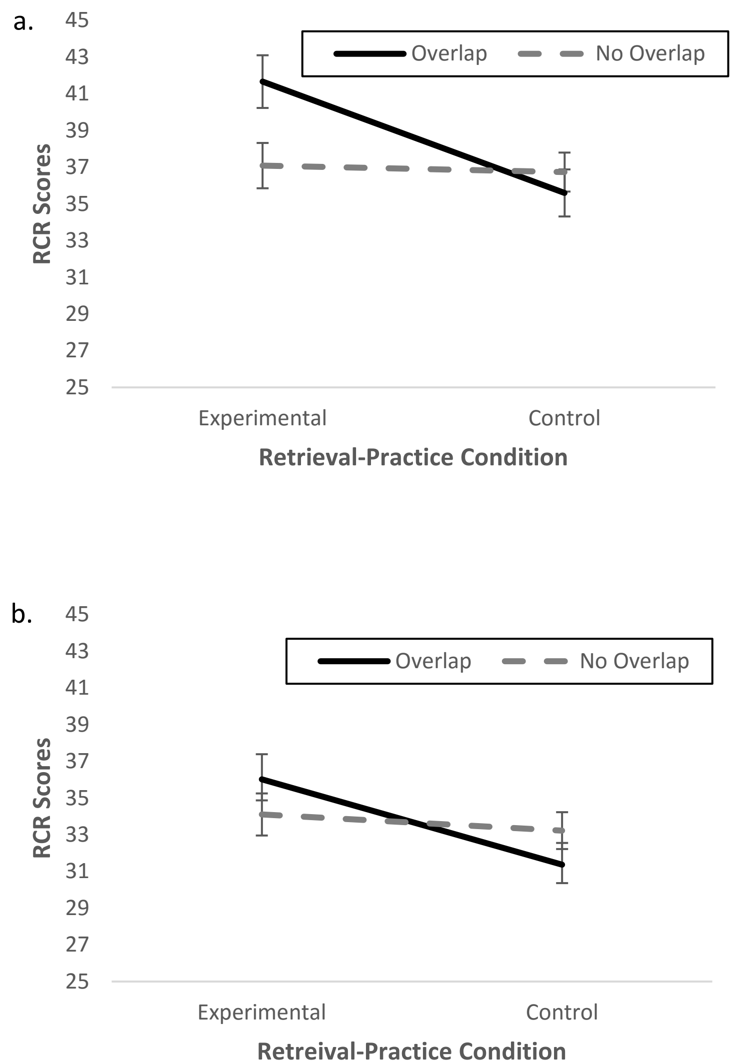

Figure 4. Performance of Ps on target propositions in the fact-verification task (a) and on the filler propositions in the fact-verification task (b). Higher RCR scores denote faster and more accurate responses.

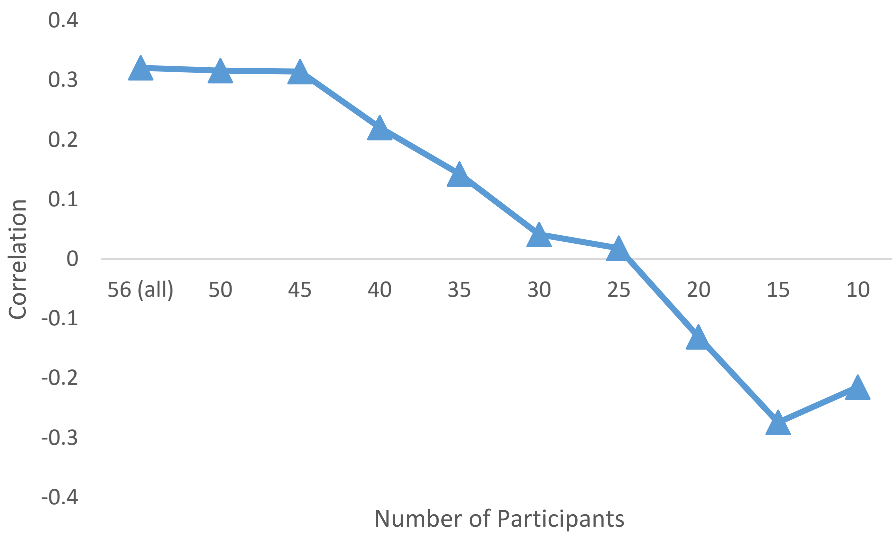

Figure 5. Plots the correlation between retrieval-practice accuracy and fact-verification performance across different group sizes. A group size of *X* consists of the *X* participants who had the highest retrieval-practice accuracy.

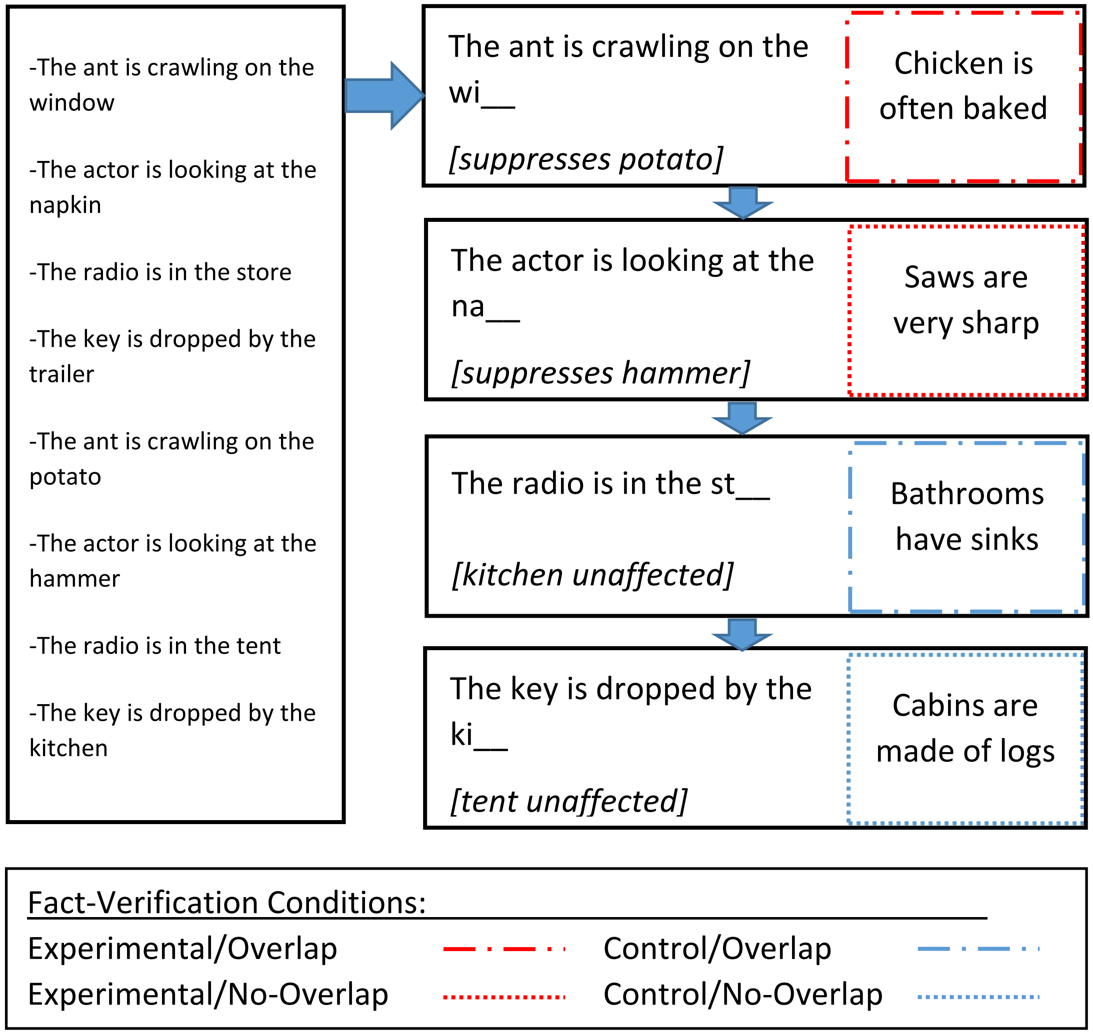

Figure 6. Design scheme for Experiment 3. After studying all of the propositions (8 listed here), Ps began the second phase, which combined the retrieval-practice task with the fact-verification task. Every retrieval-practice trial was followed by a fact-verification trial, each pair of which being termed a *unit*. The right side of each of the 4 units presented here states the fact-verification trial of that unit; the left side states the retrieval-practice trial and, underneath and in brackets, the object that has semantic overlap with the subject of the fact-verification trial along with its status (i.e., suppressed or unaffected).

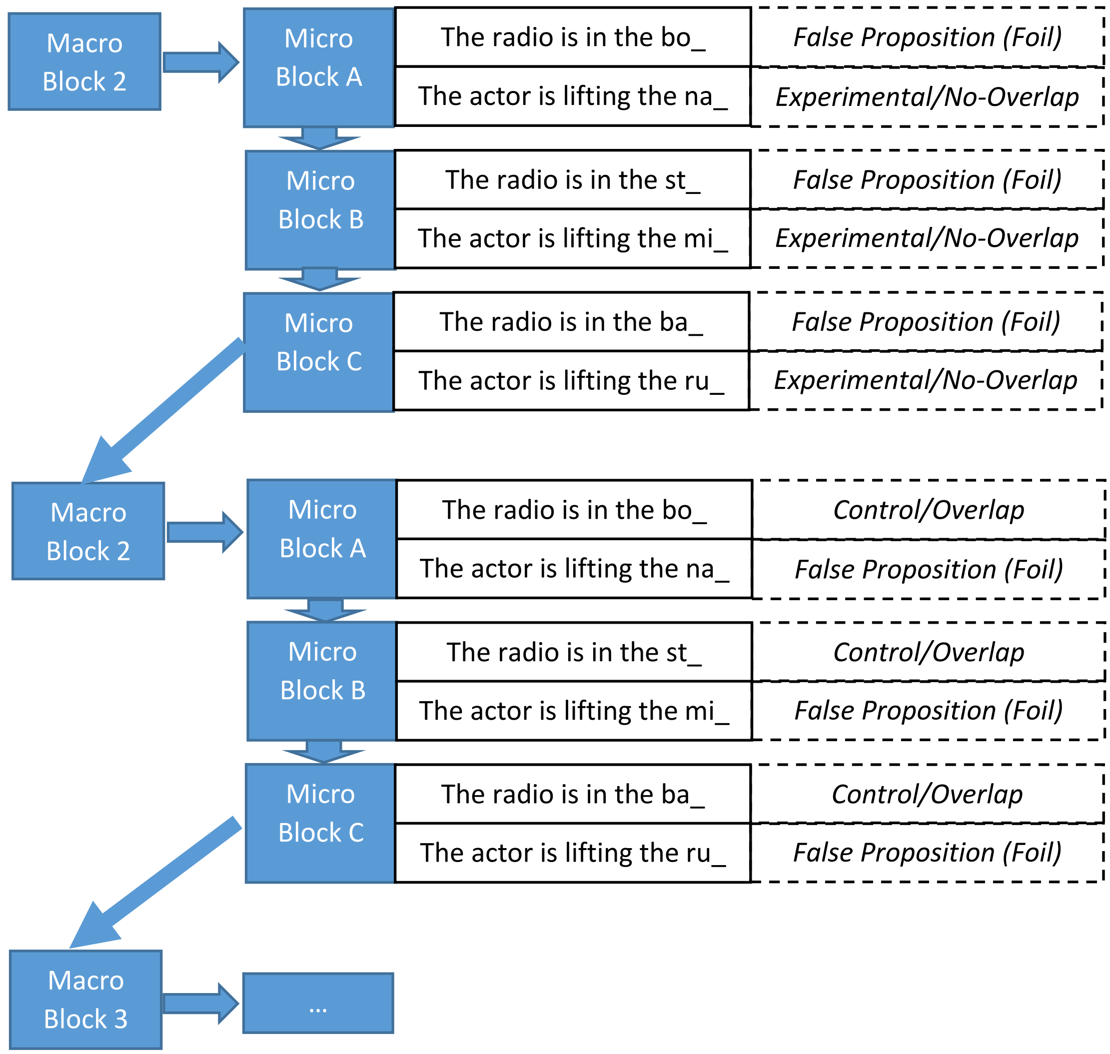

Figure 7. Design for Phase 2 of Experiment 3. The phase was split into 4 *macro blocks*, 2 of which are shown here. Each *macro block* consisted of 3 *micro blocks*, and Ps did retrieval-practice on one object for each of the six subjects in a *micro block*. Only 3 objects per subject receive retrieval-practice throughout the entirety of Phase 2, meaning that each of these objects received retrieval-practice once in each *macro block* and therefore received retrieval-practice a total of four times. In two of the *macro blocks*, the retrieval-practice of a given object was followed by a fact-verification trial over a false proposition (filler). In the other 2 *macro blocks*, retrieval-practice of that object was followed by fact-verification trail over a true proposition, once in the *experimental* condition and once in the *control* condition.

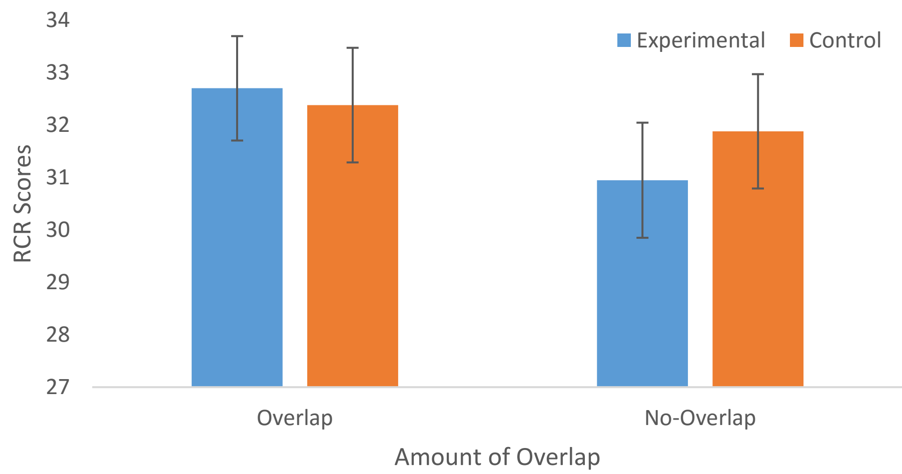

Figure 8. Performances in Experiment 3. The *pattern-suppression model* predicted that performances on *experimental* fact-verification trials would be worse than performances on *control* trials in the *overlap* condition, but that there should be no difference between the performances in the *no-overlap* condition. These predictions were not borne out.