



Semantic knowledge influences whether novel episodic associations are represented symmetrically or asymmetrically

Vencislav Popov^{1,2} · Qiong Zhang^{2,3} · Griffin E. Koch^{2,4,5} · Regina C. Calloway^{4,5} · Marc N. Coutanche^{2,4,5}

© The Psychonomic Society, Inc. 2019

Abstract

We provide new evidence concerning two opposing views of episodic associations: The independent-association hypothesis posits that associations are unidirectional and separately modifiable links ($A \rightarrow B$ and $A \leftarrow B$); in contrast, the associative-symmetry hypothesis proposes that a single, bidirectional association exists between A and B ($A \leftrightarrow B$). We used a novel method to demonstrate that whether or not episodic associations are symmetric depends on whether there is a preexisting semantic relationship between A and B. In two experiments, participants studied 30 semantically unrelated and 30 semantically related pairs intermixed in a single list and then performed a series of up to eight cued-recall test cycles. All pairs were tested in each cycle, and the testing direction ($A \rightarrow ?$ or $B \rightarrow ?$) alternated between cycles. Unrelated pairs exhibited associative symmetry—that is, accuracy and response times improved gradually on each test—suggesting that testing in both directions strengthened the same association. In contrast, semantically related pairs exhibited a stair-like pattern, in which performance did not change from odd to even tests when the test direction changed; it only improved between tests in the same direction. We concluded that episodic associations can have either a single bidirectional representation or separate directional representations, depending on the semantic relatedness of their constituent items.

Keywords Associative symmetry · Episodic memory · Cued recall · Semantic relations

Episodic and semantic associations in memory interact in various ways, and the nature of that interaction has been a matter of intense debate (Anderson & Ross, 1980; Greenberg & Verfaellie, 2010; McCloskey & Santee, 1981; McKoon, Ratcliff, & Dell, 1985, 1986). Understanding this interaction is fundamental to developing an integrated theory of memory and has spurred decades of research concerning how the presence of semantically related information affects novel episodic memory, in both positive (Ahmad, Fernandes, & Hockley,

2015; Badham, Estes, & Maylor, 2012; Epstein, Phillips, & Johnson, 1975; Guttentag, 1995; Silberman, Bentin, & Mikkulainen, 2007) and negative (Popov, Hristova, & Anders, 2017; Roediger & McDermott, 1995) ways.

At the same time, there are still gaps in our understanding of how semantic relatedness affects episodic memory. For example, consider the question of how people represent a novel episodic association between a pair of items, $A-B$. One possibility, reflected by the associative-symmetry hypothesis (ASH), is that the two items are connected in episodic memory by a single symmetric and bidirectional association ($A \leftrightarrow B$; Asch & Ebenholtz, 1962; Kahana, 2002). A second possibility, reflected by the independent-association hypothesis (IAH), is that there are separate forward ($A \rightarrow B$) and backward ($A \leftarrow B$) connections between the individual items, and that these associations can be accessed and strengthened independently from one another (Wolford, 1971). This question has been hotly debated for decades because it has implications for formal models of episodic memory, which need to explicitly specify the nature of the associative representation (Kahana, 2002; Kato & Caplan, 2017). Although recent research has for the most part supported the symmetry account (Caplan, Glaholt, & McIntosh, 2006; Kahana, 2002; Madan,

✉ Vencislav Popov
vencislav.popov@gmail.com

¹ Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, USA

² Center for Neural Basis of Cognition, Pittsburgh, PA 15213, USA

³ Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA

⁴ Department of Psychology, University of Pittsburgh, Pittsburgh, PA 15260, USA

⁵ Learning Research & Development Center, University of Pittsburgh, Pittsburgh, PA 15260, USA

Glaholt, & Caplan, 2010; Sommer, Rose, & Büchel, 2007; Sommer, Schoell, & Büchel, 2008) for associations between unrelated items (e.g., BAKER–MOUNTAIN), to our knowledge there has been almost no research on whether the directionality of episodic associations is affected by the presence of a semantic relationship between an association's components (e.g., HUNTER–FOREST). The goal of the present article is to address this question.

Associative symmetry

Associative symmetry has been studied with two main paradigms. The initial research compared overall memory for forward (A–?) and backward (B–?) cued-recall tests. In support of the ASH, most studies have shown that forward and backward recall levels were equivalent when the cue and response were drawn from the same population of words (Asch & Ebenholtz, 1962; Ekstrand, 1966; Horowitz, Norman, & Day, 1966; Murdock, 1966; for a review, see Kahana, 2002). However, a number of researchers have pointed out that these results do not necessarily support the ASH. For example, Wolford (1971) argued that equal forward and backward recall could occur not only if there were a single, bidirectional association, but also if two distinct associations (A to B and B to A) were formed with the same strength (Ebbinghaus, 1885/1913; Robinson, 1932). This argument was later reiterated by Kahana through a formal analysis of existing mathematical models of associative memory. Similarly, others have argued that the few cases in which an asymmetry has emerged between forward and backward recall (Bartling & Thompson, 1977; Lockhart, 1969) are not necessarily inconsistent with symmetric models—asymmetry could occur not only because of an asymmetric representation, but also if the two items differed in availability (Asch & Ebenholtz, 1962; Horowitz et al., 1966) or in their numbers of preexisting associations (Kahana, 2002).

Because of these interpretative complications, Kahana (2002) suggested that the comparison of forward and backward recall cannot discriminate between the ASH and IAH. Instead, he argued that this question could be addressed by calculating the correlation in recall between successive forward and backward tests at the level of individual pairs. On the basis of a formal model analysis, Kahana reasoned that if associations are stored symmetrically, the correlations in recall should not differ depending on whether both tests are performed in the same direction (i.e., both A–? or both B–?) or in different directions (one A–? and one B–?). Kahana found near perfect correlation between same and different directions, a result that has been replicated multiple times for word pairs (Caplan et al., 2006; Kahana, 2002; Madan et al., 2010; Rehani & Caplan, 2011) and for object–location associations (Sommer et al., 2007; Sommer et al., 2008). This correlation

in performance is not affected by the number of study repetitions (Kahana, 2002) or by some stimulus properties, such as word frequency or imageability (Madan et al., 2010). These results have led to a tentative consensus in favor of the ASH, although no current model can easily account for the full pattern of results concerning the representation of order and direction (Kato & Caplan, 2017; Rehani & Caplan, 2011).

The effect of semantic relatedness

Given the considerable interest in how semantic and episodic memory interact, it is surprising that there has been almost no research on the impact of semantic relatedness on associative symmetry (one notable exception is Caplan, Boulton, & Gagné, 2014, which we discuss below). Whereas a large proportion of lab-based research has been based on learning arbitrary novel associations between random combinations of items, arguably, much real-world learning involves linking meaningfully related pieces of information. Consequently, the lack of data on this question impedes progress in developing integrated models of semantic and episodic memory.

Why should we expect semantic relatedness to impact associative symmetry? In contrast to unrelated word pairs (e.g., BAKER–MOUNTAIN), semantically related word pairs (e.g., BELLHOP–HOTEL) have preexisting associations in long-term memory (LTM). Pairs of semantically related words are not only easier to remember (Epstein et al., 1975; Guttentag, 1995; Silberman et al., 2007), but their storage also involves different underlying learning mechanisms. The beneficial effects of semantic relatedness have been attributed to a process of rapid integration into cortical networks that might bypass hippocampal learning (Bein, Reggev, & Maril, 2014; Coutanche & Thompson-Schill, 2015; McClelland, 2013; Tse et al., 2007; van Kesteren, Ruitter, Fernández, & Henson, 2012), to a boost from preexisting associative strength (Silberman et al., 2007), to unitization (Parks & Yonelinas, 2015), or to the encoding of the abstract relation between the words (Popov et al., 2017). As we explain below, any one of these mechanisms might lead to the formation of directional episodic associations between semantically related items.

First, as revealed by free association norms, the strength of semantic associations are often asymmetrical (Nelson, McEvoy, & Schreiber, 2004). For example, BELLHOP is very strongly related to HOTEL, while HOTEL is only weakly related to BELLHOP—people are much more likely to generate HOTEL in response to BELLHOP, but not vice versa. Second, even when the strength of the association is equal in both directions, the meaning of how A is related to B is often different from the meaning of how B is related to A. For example, a FOREST HUNTER is a hunter who hunts in forests, whereas a HUNTER FOREST is a forest meant for hunting (Coutanche, Solomon, & Thompson-Schill, *in press*).

In a similar argument, Kahana (2002) suggested that differences in forward and backward recall can occur if the cue changes the interpretation of the target, resulting in two different representations, A–B' and B–A'. Even though Kahana made this argument with respect to random pairings of words, we believe it applies to an even greater extent to semantically related pairs.

A number of empirical findings are consistent with this possibility. For example, when participants study a word pair such as COOKIE JAR, and are later asked to recognize the modifier word (COOKIE), they do better when it is embedded in a relationally similar word pair (COOKIE plate) than when it is in a relationally dissimilar pair (COOKIE crumb; Jones, Estes, & Marsh, 2008). Further support comes from studies on *unitization*—when two words can form a compound because they are semantically related, people tend to unitize them and represent them as a single item, rather than as an association between two items. This results in more familiarity-based responses (Ahmad et al., 2015), a reduction in binding related age impairments (Ahmad et al., 2015; Badham et al., 2012) and increased neural markers of familiarity in ERP (Kriukova, Bridger, & Mecklinger, 2013; Rhodes & Donaldson, 2007) and fMRI data (Ford, Verfaellie, & Giovanello, 2010; Haskins, Yonelinas, Quamme, & Ranganath, 2008). Thus, people might store semantically related word pairs such as HUNTER FOREST and FOREST HUNTER as different representations, even though the constituent elements are the same.

Furthermore, when presented with semantically related words, people automatically activate and encode the specific semantic relation present in the stimuli, regardless of whether or not that is required by the task (Popov & Hristova, 2015; Popov et al., 2017). Popov et al. (2017) found that when participants study word pairs (e.g., FLOOR–CARPET) and then have to distinguish between intact and recombined pairs, they make more false alarms and respond more slowly to lures that have the same semantic relation as some of the studied pairs (e.g., the lure TABLE–CLOTH shares the *X is covered by Y* relation with the studied pair FLOOR–CARPET). Thus, people store not only the individual items, but also the semantic relation between them. In contrast to episodic associations, for which directionality is often of no import, the order of items in most semantic associations determines the meaning of the relation—for example, *John loves Mary* and *Mary loves John* represent different states that are not necessarily reciprocal. Thus, even if people do not unitize semantically related pairs, as was suggested in the previous paragraph, they might store different relations for the pair SCHOOL–TEACHER and the pair TEACHER–SCHOOL: either *X-is-a-workplace-for-Y* or *X-works-in-Y*.

Preliminary evidence for associative asymmetry for semantically related items has come from Caplan et al. (2014), who tested forward and backward recall on two

successive tests for unrelated words (BRIEF–TAX) and compound words whose meaning was either transparent from the meaning of their constituents (ROSE–BUD), or not (HOG–WASH). Replicating previous results, for unrelated words they found a symmetric correlation between successive tests in the same or different directions. However, all compound word conditions showed associative asymmetry with reduced test–retest correlation when the direction between tests changed. The interpretation of these promising results, however, is complicated by the fact that ROSE–BUD and HOG–WASH form existing words, whereas BUD–ROSE and WASH–HOG do not. Thus, the results might be due to the fact that there is no need to form an episodic association between ROSE and BUD—instead, one can store the existing item. A further issue, as we discuss below, is that finding a reduced correlation between successive tests is not necessarily an indicator of associative asymmetry in the representation.

Present experiments

On the basis of the preceding discussion, we expected that semantically related word pairs might not show the associative symmetry usually seen with semantically unrelated pairs. Although the currently most popular method for testing associative symmetry involves calculating the correlation between two successive tests of the same pair, recent research has suggested that a reduced correlation does not necessarily reflect a violation of associative symmetry; factors at retrieval can also reduce the correlation (Rehani & Caplan, 2011). Rehani and Caplan used the successive-testing procedure on two types of stimuli—single function lists, in which each word appeared in only one pair, and double-function lists, in which each word appeared in two different pairs. They found that the correlation between tests was smaller for double-function lists, *even between tests in the same direction*, suggesting that interference during retrieval could result in a reduced test–retest correlation.

To assess associative symmetry in this investigation, we examined the learning curve for individual pairs as a result of multiple intermixed forward (A–?) and backward (B–?) cued-recall tests (i.e., A–?, B–?, A–?, B–?, etc.). Testing, especially if followed by a restudy opportunity, improves memory (Liu, Liang, Li, & Reder, 2014; Roediger & Karpicke, 2006; Rowland, 2014). This improvement possibly results from strengthening the association between the cue and the target on each successive test. We reasoned that if an association is represented symmetrically and both directions are strengthened at once, then performance as measured by recall accuracy and response times should increase monotonically from

one test to the next, regardless of the fact that tests alternate in direction. In contrast, if an association is represented as two separate links that can be modified independently, then testing on one cycle would strengthen only the cue-to-target association, but not the target-to-cue association. This should result in a stair-like pattern, in which performance does not improve from Test 1 to Test 2 (or odd to even tests in general), because Test 1 strengthens only the $A \rightarrow B$ association, and Test 2 requires retrieving the $A \leftarrow B$ association instead. Performance in this case would improve only from Test 1 to 3 to 5 to 7, and from Test 2 to 4 to 6 to 8, because those tests involve the same direction. According to previous research, semantically unrelated pairs (e.g., BAKER–MOUNTAIN) are represented symmetrically (Caplan et al., 2006; Kahana, 2002; Madan et al., 2010; Rehani & Caplan, 2011); thus, we predicted that these word pairs would show a gradual increase in performance over test cycles. In contrast, the shape of the learning curve for semantically related pairs (e.g., TEACHER–SCHOOL) is an open question. We report results from two experiments in which we manipulated the semantic relatedness of the words in each pair and examined whether there would be transfer in learning from one test to another when the tests changed directions.

Experiment 1

The first experiment was not specifically designed to test the questions raised in this article (see Zhang, Popov, Koch, Calloway, & Coutanche, 2018, for the original purpose of this experiment). Nevertheless, we realized that, fortunately, the design of the repeated cued-recall testing task allowed us to test our predictions. Participants performed another task after they completed the procedure described below. Since that task is irrelevant for the present analysis and was performed after the task described below, we have omitted it from the description (a description is available in Zhang et al., 2018). All materials, experimental software, data, and analysis code are available at <https://osf.io/72amw>.

Method

Participants Thirty-five members of the Pittsburgh community participated in the study in exchange for either course credit or payment (18 males, 17 females; $M_{\text{age}} = 20.7$ years, $SD = 3.0$; all were English speakers without a learning or attentional disorder). The University of Pittsburgh Institutional Review Board approved all procedures. Three participants were excluded from the analysis because they did not reach criterion for at least half of the studied pairs (see the Procedure section).

These exclusion criteria were established prior to the start of data collection.

Materials Eighty words representing professions/people and locations were split into 20 sets of four words each (see Table 1). Half of the sets contained semantically unrelated words, and the other half contained words that were pairwise semantically related. The four words in each set (A, B, C, and D) formed three word pairs, A–B, B–C, and C–D. Each pair always consisted of one word representing a person and one word representing a location. The resulting 60 word pairs were studied in the experiment.

The semantically related words were selected on the basis of latent semantic analysis (LSA), a method that approximates global semantic knowledge through the local co-occurrence of words in subsamples of language (Landauer & Dumais, 1997). LSA scores range from -1 to 1 , with higher values indicating stronger semantic associations. LSA does not distinguish the order or the type of semantic relations. Word pairs in the semantically related condition had higher LSA scores ($M_{\text{LSA}} = .42$, $SE = .04$) than word pairs in the semantically unrelated condition ($M_{\text{LSA}} = .06$, $SE = .01$) [$t(58) = 8.59$, $p < .001$]. The words in each condition did not differ in length ($M_{\text{rel}} = 6.62$, $M_{\text{unrel}} = 6.69$) [$t(79) = 0.17$, $p = .86$] or normative frequency ($M_{\text{rel}} = 39.6$ words per million, $M_{\text{unrel}} = 41.2$ words per million) [$t(79) = 0.11$, $p = .91$].

Procedure All Experimental procedures were created and presented using the PsychoPy 2 Experiment Builder, version 1.84.2 (Peirce, 2007). The experiment had two phases—an initial paired-associate learning task and repeated cued-recall tests, which were performed to a criterion and involved a restudy of the materials.

During the initial study phase, word pairs were presented one at a time in the middle of the screen for 3.5 s., separated by a fixation cross for 0.5 s (see Fig. 1). The two words were presented one above the other. Participants had to remember the word pairs for a subsequent cued-recall test. Their task during the study phase was to rate how likely it would be to see the person in the corresponding location, using a 4-point scale (*very unlikely* to *very likely*).

After all 60 pairs had been studied, participants performed a series of cued-recall test cycles. In each cycle, all 60 pairs were tested and restudied in a new random order. On each test trial, participants saw one word as a cue, and they had 5 s to type in the target word that had been associated with the cue during study. If participants did not response within 5 s, the trial was considered incorrect. Because some words were studied in more than one pair, the first letter of the correct target was also shown. After participants had responded, the correct answer was displayed along with the cue for restudy. When the response was incorrect, the cue word was tested again *in this cycle*, following the presentation of all other

Table 1 Stimuli used in Experiment 1

Condition	Stimuli			
	A	B	C	D
Semantically related	Teacher	Classroom	Student	Dormitory
	Doctor	Hospital	Patient	Reception
	Judge	Courtroom	Lawyer	Firm
	Astronaut	Spacecraft	Pilot	Cabin
	Fireman	Station	Policeman	Precinct
	Hunter	Forest	Ranger	Park
	Chef	Kitchen	Waiter	Parlor
	Player	Arena	Announcer	Newsroom
	Bellhop	Hotel	Guest	Mansion
	Driver	Road	Racer	Stadium
Semantically unrelated	Baker	Mountain	Principal	Circus
	Scientist	Street	Nurse	Tavern
	Cyclist	Inn	Prosecutor	Library
	Maid	Track	Explorer	Casino
	Surgeon	School	Officer	Lodge
	Guard	Restaurant	Chief	Ballroom
	Professor	Rocket	Hostess	City
	Gladiator	Chamber	Coach	Ward
	Lumberjack	Highway	Clerk	Airport
	Bailiff	Pharmacy	Motorist	Office

Each row represents one set of items. Participants studied word pairs formed from each set—that is, AB, BC, and CD

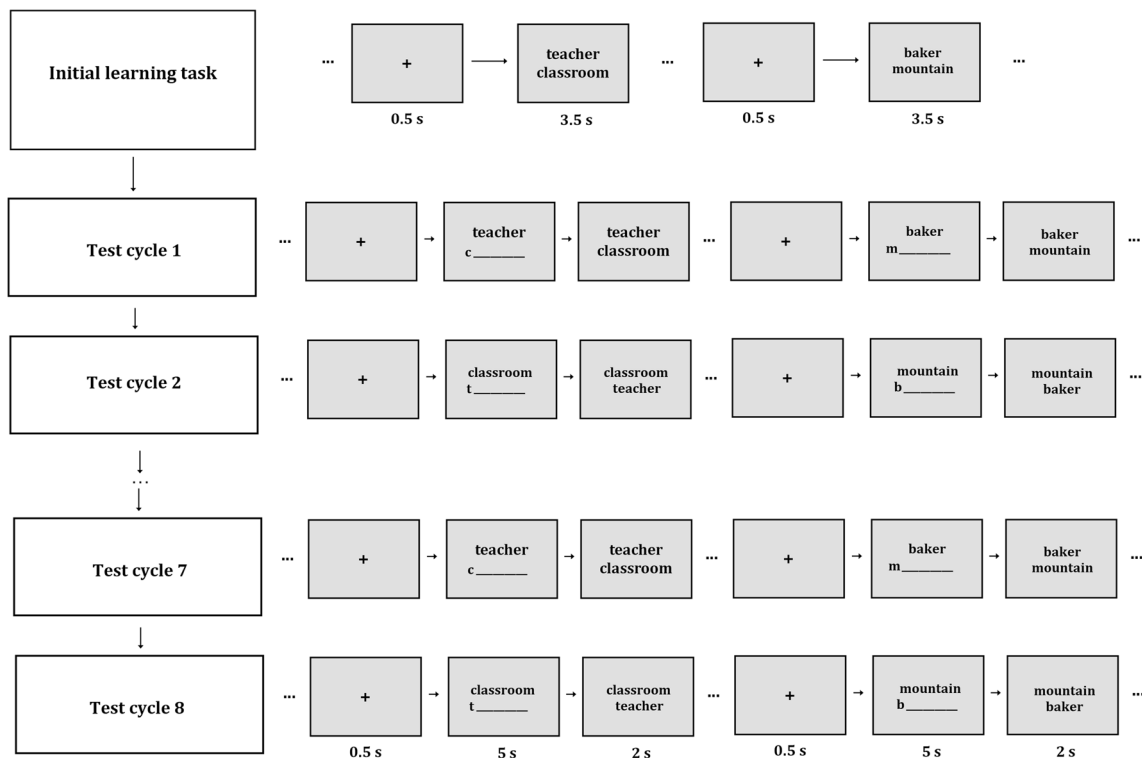


Fig. 1 Illustration of the experimental procedure. Each row is a different phase, showing two example trials within that phase for a semantically related and a semantically unrelated word pair

cue words. This procedure continued until all pairs had been correctly recalled, which concluded one cycle.

In the first cycle, all pairs were tested in the forward direction (e.g., A-?). The subsequent cued-recall cycles were identical to the first, but the test order of each pair was reversed—pairs that were tested in the forward direction on one cycle were then tested in the backward direction on the subsequent cycle, and vice versa (see Fig. 1 for an illustration of the task sequence). The task ended after three consecutive cycles with perfect recall of all pairs. Participants performed between 6 and 12 cycles before reaching this criterion for all pairs. Of the 32 participants, 30 completed a minimum of eight cycles (two participants reached criterion in only six cycles), and only 18 of those 30 completed more than eight cycles. Because of this, we limited our analyses to Cycles 1–8.

Design This was a within-subjects experiment with two factors—semantic relatedness (related vs. unrelated) of the words in each pair, and the test number (1 to 8). The order of the words in each pair and the order of pairs in the study list were randomized for each participant, and each test was presented in a new random order for each participant.

Data analysis We compared the first keypress response times (RTs) and cued-recall accuracy on odd versus even test cycles. For each trial, the cued-recall accuracy was scored as 1 when the participant typed the correct target associated with the cue, and as 0 when the participant either provided a wrong response or did not respond at all. Misspelled (correct) responses were counted as correct. Because the odd and even test cycles tested memory for the word pair from two different directions (A-? vs. B-?), improvement from odd to even tests would reflect a bidirectional strengthening of the association. We used logistic (for accuracy) and linear (for RTs) mixed-effect regression analyses (Baayen, Davidson, & Bates, 2008), as implemented in the lme4 package (Bates, Mächler, Bolker, & Walker, 2015) in R (R Core Team, 2014). The key question for the analysis was how learning would transfer from one test to another when the test order changed direction—that is, to determine whether performance would improve gradually or instead would follow a step-like pattern. Since the tests were always A-? on odd cycles and B-? on even cycles, one reasonable analysis was to test whether performance improved from odd to even cycles and whether that improvement interacted with the semantic relatedness factor. A decrease in RTs or an increase in accuracy from odd to even cycles would signify that testing on A-? led to improvement on the following B-? test cycle. Thus, the regression models included fixed effects for semantic relatedness, odd versus even test cycles, and the interaction between the two. The regression models also included random by-subject intercepts and random by-

subject slopes for the semantic relatedness and odd-versus-even factors, as well as their interaction. We inferred the significance of each effect on the basis of likelihood ratio tests and Akaike/Bayesian information criterion (AIC/BIC) comparisons of the regression model that contained the effect in question with the identical model that lacked this contrast (a decrease in AIC/BIC greater than 2 is conventionally considered to provide evidence for a better fit). For completeness, we also report the corresponding Bayes factors (BFs), calculated with the BIC approximation of the BF (Wagenmakers, 2007). BF_{null} reports evidence in favor of the null model, and BF_{alt} reports evidence in favor of the alternative model that contained the effect of interest. The RT analyses were performed only on correctly recalled items.

Results and discussion

Figure 2 shows the RTs (left) and cued-recall accuracy (right) on each test cycle separately for semantically related and semantically unrelated word pairs. Performance for the semantically unrelated word pairs improved on each test cycle—RTs decreased and accuracy increased gradually on each test. In contrast, performance for the semantically related pairs improved only between odd cycles that tested memory in the same direction (A-?) and between even cycles that tested memory in the opposite direction (B-?). Critically, there was no improvement in either RTs or accuracy from odd to even tests, as evidenced by the stair-like pattern in Fig. 2. The mixed-effect models revealed significant main effects of semantic relatedness on both RTs, $\Delta AIC = -48$, $\Delta BIC = -41$, $\chi^2(1) = 50.26$, $p < .001$, $BF_{\text{alt}} > 150$, and accuracy, $\Delta AIC = -65$, $\Delta BIC = -57$, $\chi^2(1) = 66.68$, $p < .001$, $BF_{\text{alt}} > 150$; no main effect of even-versus-odd tests on RTs, $\Delta AIC = 2$, $\Delta BIC = 15$, $\chi^2(1) = 0.10$, $p = .75$, $BF_{\text{null}} > 150$, but a significant main effect of even-versus-odd tests on accuracy, $\Delta AIC = -27$, $\Delta BIC = -19$, $\chi^2(1) = 29.37$, $p < .001$, $BF_{\text{alt}} > 150$; and a significant interaction between semantic relatedness and even-versus-odd tests for both RTs, $\Delta AIC = -10$, $\Delta BIC = -3$, $\chi^2(1) = 11.77$, $p < .001$, $BF_{\text{alt}} = 4.5$, and accuracy, $\Delta AIC = -5$, $\Delta BIC = -2$, $\chi^2(1) = 7.45$, $p = .006$, $BF_{\text{alt}} = 2.7$. Post-hoc tests with the Bonferroni–Holm correction for multiple comparisons revealed that the improvement from odd to even tests occurred only for semantically unrelated pairs (for RTs, $\Delta = -70$ ms, $z = -2.55$, $p = .021$; for accuracy, $\Delta = 10\%$, $z = 8.131$, $p < .001$), but not for semantically related pairs (for RTs, $\Delta = 40$ ms, $z = 1.80$, $p = .135$; for accuracy $\Delta = 3\%$, $z = 1.707$, $p = .166$). In summary, semantically unrelated pairs exhibited bidirectional strengthening, whereas the pattern for semantically related pairs was more consistent with a model in which the forward and backward associations were strengthened independently from one another.

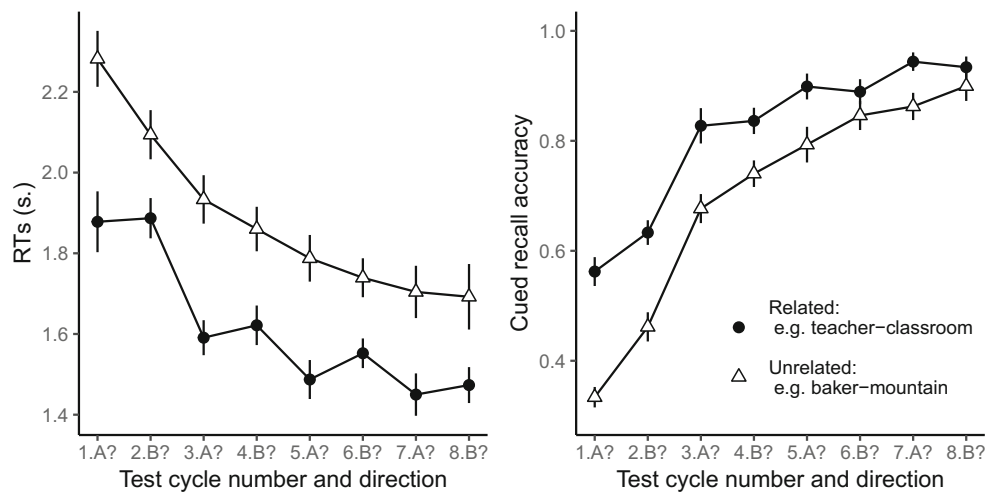


Fig. 2 Cued-recall response times (RTs, left) and accuracy (right) in Experiment 1, depending on the semantic relatedness of the words in each pair and the test number/direction. Pairs were studied originally in

the AB order and were tested in either the forward (A-?) or the backward (B-?) direction (in each cycle). The test direction alternated on subsequent tests. Error bars represent 95% confidence intervals

A few aspects of Experiment 1's design complicate the interpretation of the results. First, all participants studied the same sets of semantically related and unrelated pairs, and these pairs were designed to contain one word representing a person/profession and one word representing a location. Although this does not negate the findings for the set of stimuli we used, the results might be specific to these person–location stimuli, and it is unknown whether all semantically related pairs would show associative asymmetry. Another issue is that we used double-function lists (Rehani & Caplan, 2011; Slamecka, 1976), in which words were part of more than one word pair (i.e., AB, BC, and CD). It is possible that the stair-like pattern for semantically related words was the result of interference from the shared cues. Consider a pair AB, which is tested with the A cue on Cycle 1 (A-?). Then, on Cycle 2, the same pair is tested with the B cue (B-?). However, the B cue could also lead to retrieval of the BC pair, which could cause interference and reduce performance. Thus, even if performance would have been boosted from Cycle 1 to Cycle 2 and from Cycle 3 to Cycle 4, and so forth, this boost might have been masked by interference. Why would interference occur only for semantically related, but not for unrelated pairs? As we demonstrated, memory for semantically related pairs was significantly higher overall, which means that when cueing AB with B, the representation of the competing pair, BC, would also be stronger, potentially resulting in more interference. We note, however, that if this explanation were correct, we should have observed only a quantitative difference between related and unrelated pairs, rather than the qualitative shift from a stair-like to a smooth form of the pattern we obtained. There was no evidence for any interference for unrelated pairs, which showed a smooth improvement in performance (RTs and accuracy) over test cycles. Furthermore, it is unclear why interference would

occur only on even but not on odd tests, resulting in the stair-like pattern of results for semantically related pairs. Nevertheless, a cleaner design in which the words in each pair were unique would more convincingly discount this possible explanation. The goal of Experiment 2 was to replicate the results of Experiment 1 using the two above-mentioned controls.

Experiment 2

We used a broader range of semantic relations drawn from Popov et al.'s (2017) pool of 106 word pairs, which are exemplars of 30 semantic relations. Each participant saw a random sample of 30 of those pairs in the semantically related condition, and a random sample of 30 pairs was recombined to form the pairs for the semantically unrelated condition. All words appeared in only one word pair, removing the potential interference from using double-function lists. The goals, methods, and main analysis of Experiment 2 were preregistered at <https://osf.io/rdsw5/>. All materials, experimental software, and data are available at <https://osf.io/72amw>.

Method

Participants We aimed to collect data from 30 participants. Thirty-five University of Pittsburgh students participated in exchange for course credit or payment. Five of the participants were excluded from the analysis because they did not complete all six cycles within the allotted time. These participants were replaced in order to obtain the goal sample of 30 people. The University of Pittsburgh Institutional Review Board approved all procedures.

Materials All words were drawn from Popov et al.'s (2017) pool of 106 word pairs, which are exemplars of 30 semantic relations.¹ LSA scores for each pair were calculated with the LSAfun package (Günther, Dudschig, & Kaup, 2015) for the R statistical software, and the scores were based on the EN_100k_lsa corpus, available from <http://www.lingexp.uni-tuebingen.de/z2/LSAspaces/>. For each participant, 30 word pairs (one from each relation) were selected at random to be presented in the semantically related condition. These pairs had a mean LSA score of .53, ranging from .20 to .95. The remaining 76 word pairs were recombined at random for each participant, and the 30 with the lowest LSA scores were selected for the semantically unrelated condition (mean LSA = .059, ranging from .003 to .094). In addition, none of the pairs in the semantically unrelated condition appeared in the USF Free Association Norms (Nelson et al., 2004).² The words in each condition did not differ in either length [$M_{rel} = 5.62$, $M_{unrel} = 5.67$, $t(413) = 0.24$, $p = .81$] or normative frequency [$M_{rel} = 50.8$, $M_{unrel} = 49.3$ words per million; $t(79) = 0.17$, $p = .87$]. Due to the randomization of stimuli per participant, each of the 106 related word pairs was presented to a median of eight participants (ranging from 1 to 13, $SD = 2.24$), and each of the 853 unrelated word pairs was presented to a median of one participant (ranging from 1 to 3, $SD = 0.209$).

Procedure and design The procedure was identical to the one used in Experiment 1, with several exceptions: (1) Each studied word pair contained unique words not shared with any other pair; (2) since each cue was associated with only one target, the cued-recall test did not show the first letter of the target; (3) all participants had to complete exactly six cued-recall cycles within 1 h; and (4) during the study phase, participants rated how semantically related the items were, on a

scale from 1 (*not at all related*), through 2 (*not related*) and 3 (*somewhat related*), to 4 (*strongly related*).

Data analysis The data analysis procedure was identical to that in Experiment 1. Two of the semantically related word pairs were consistently rated by participants in the study phase as being unrelated (relatedness scores of 1.33 and 2.00). We excluded these pairs from the analysis.

Results

We replicated the main pattern of results (see Figs. 3 and 4) from Experiment 1. The cued-recall performance for semantically unrelated pairs improved gradually on each test cycle, regardless of test direction. In contrast, the performance for semantically related pairs again showed the stair-like pattern we had observed in Experiment 1, although it was less pronounced. The mixed-effect models revealed significant main effects of semantic relatedness on both RTs, $\Delta AIC = -43$, $\Delta BIC = -37$, $\chi^2(1) = 45.393$, $p < .001$, $BF_{alt} > 150$, and accuracy, $\Delta AIC = -41$, $\Delta BIC = -38$, $\chi^2(1) = 43.283$, $p < .001$, $BF_{alt} > 150$; a main effect of even-versus-odd tests on RTs, $\Delta AIC = -9$, $\Delta BIC = -1$, $\chi^2(1) = 10.550$, $p = .001$, $BF_{alt} = 1.6$; and a significant main effect of even-versus-odd tests on accuracy, $\Delta AIC = -27$, $\Delta BIC = -20$, $\chi^2(1) = 29.25$, $p < .001$, $BF_{alt} > 150$. Importantly, we also found significant interactions between semantic relatedness and even-versus-odd tests for both RTs, $\Delta AIC = -7$, $\Delta BIC = -3$, $\chi^2(1) = 8.384$, $p = .004$, $BF_{alt} = 4.5$, and accuracy, $\Delta AIC = -14$, $\Delta BIC = -5$, $\chi^2(1) = 15.776$, $p < .001$, $BF_{alt} = 12$. Post-hoc tests with the Bonferroni–Holm correction for multiple comparisons revealed that improvements in RTs and accuracy from odd to even tests occurred only for semantically unrelated pairs (for RTs, $\Delta = -118$ ms, $z = -4.224$, $p < .001$; for accuracy, $\Delta = 13\%$, $z = 8.131$, $p < .001$), but not for semantically related pairs (for RTs, $\Delta = -31$ ms, $z = -1.316$, $p = .328$; for accuracy $\Delta = 2\%$, $z = 1.769$, $p = .142$). Figure 4 shows that in both experiments, there was improvement from odd to even tests in both RTs and accuracy for unrelated, but not for related, pairs.

To test whether overall differences in accuracy between the semantically related and unrelated conditions could account for the odd–even difference in the two conditions, we calculated the average accuracy for each subject and relatedness condition, and repeated all model comparisons after including this factor in the regression models. Although, unsurprisingly, performance on each trial was successfully predicted by the overall accuracy for each condition and subject [for RTs: $\Delta AIC = -24$, $\Delta BIC = -16$, $\chi^2(1) = 25.655$, $p < .001$, $BF_{alt} > 150$; for accuracy: $\Delta AIC = -83$, $\Delta BIC = -77$, $\chi^2(1) = 85.274$, $p < .001$, $BF_{alt} > 150$], including this factor in the models did not change the substantive conclusions or the

¹ The words used by Popov et al. (2017) are in Bulgarian, and the original pool consists of 165 word pairs organized into 30 relations. After translation, 50 word pairs had to be excluded for a variety of reasons—repeated words across pairs, lack of an English single-word noun to represent some of the words, and so forth. An additional nine word pairs were excluded because their LSA scores were less than .20.

² As is noted in our preregistration materials (<https://osf.io/rdsuw5/>), we were planning to analyze associative asymmetry at the item level as a function of asymmetry in free association norms. Unfortunately, we could not perform this analysis for a couple of reasons: First, very few of the 212 words used in the study appeared as cues in the Nelson et al. (2004) association norms, preventing us from calculating the relevant predictors. A recent study has expanded these association norms to 12,000 different cues (De Deyne, Navarro, Perfors, Brysbaert, & Storms, 2018). In all, 72% of the words used in our study were presented as cues in De Deyne's dataset, but we needed norms for both words in each pair in order to calculate preexisting associative asymmetry. Of the 106 word pairs, only 51 had associative ratings for both words. The mean association strength for these 51 pairs was 6.2% ($SD = 6.7\%$), and the mean difference between forward and backward association strength was 7.4% ($SD = 7.3\%$). In addition, due to the randomization of stimuli per participant, each word pair was studied by a median of only eight participants, which means that an item-level analysis would be based on an average of only eight observations per pair.




Fig. 3 Cued-recall response times (RTs, left) and accuracy (right) in Experiment 2, depending on the semantic relatedness of the words in each pair and the test number/direction. The pairs were studied originally

in the AB order and then tested in the forward (A-?) or backward (B-?) direction (alternating with each cycle). Error bars represent 95% confidence intervals

levels of significance for any of the model comparisons presented above. Thus, the overall difference in accuracy between the semantically related and unrelated conditions cannot account for our results.

Accuracy was at ceiling from Cycle 3 onward for semantically related pairs, and since it could be argued that the comparison above is biased because there cannot be improvement from ceiling, we also compared whether accuracy improved




Fig. 4 Comparison of the Experiment 1 (left) and Experiment 2 (right) results. The top panels show RTs, and the bottom panels show cued-recall accuracy. Data are collapsed over the types of pair (semantically unrelated vs. semantically related) and whether the test was odd or even. Both experiments show improvement from odd to even tests in both RTs and

accuracy for unrelated, but not for related, pairs. Error bars represent 95% confidence intervals. Note that the y-axis scales have been adjusted to show the pattern similarity, despite differences in overall memory between the experiments

from Test Cycle 1 (A-?) to 2 (B-?), when accuracy was lower.³ As with the previous analysis, we observed a significant interaction between test cycle and semantic relatedness, $\Delta\text{AIC} = -20$, $\Delta\text{BIC} = -18$, $\chi^2(1) = 21.521$, $p < .001$, $\text{BF}_{\text{alt}} > 150$ —whereas accuracy improved drastically from Test 1 to Test 2 for semantically unrelated pairs ($\Delta = 24\%$, $z = 8.031$, $p < .001$), it did not change for semantically related pairs ($\Delta = 3\%$, $z = 0.943$, $p = .559$). This was not due to a ceiling effect—in contrast, accuracy improved from Test 2 to Test 3 for both semantically unrelated pairs ($\Delta = 8\%$, $z = 5.160$, $p < .001$) and semantically related pairs ($\Delta = 9\%$, $z = 6.121$, $p < .001$).

One of the goals of this experiment was to test whether the effects would generalize to a wide variety of semantic relations. For that reason, we wanted to examine whether the stair-like pattern was driven by a small subset of word pairs or whether the lack of improvement from odd to even tests applies to most semantically related pairs we used.⁴ For each semantically related pair, we calculated the average RTs and accuracy in Tests 1 and 2, and the difference between them. A difference score of 0 would mean that the pair showed associative asymmetry, whereas a difference score greater than 0 (for accuracy) or less than 0 (for RTs) would show symmetry. Figure 5 plots the distributions of difference scores for RTs and accuracy and compares them to the average difference scores for the semantically unrelated pairs (red vertical lines). Accuracy for 93% of all semantically related pairs improved less from Test 1 to Test 2 than the average improvement for semantically unrelated pairs; RTs for 87% of semantically related pairs improved less from Test 1 to Test 2 than the average improvement for unrelated pairs. A Bayesian t test (with the BIC approximation of BFs; Wagenmakers, 2007) showed that 61% of the semantically related pairs had a $\text{BF} > 3$ in favor of the null hypothesis that there was no improvement from Test 1 to Test 2, which is considered moderate evidence for the absence of an effect. Only 10% of the semantically related pairs had a $\text{BF} > 3$ in favor of the alternative hypothesis that there was an improvement from Test 1 to Test 2.

In summary, even when using single-function lists with a wide variety of semantic relations, semantically unrelated pairs exhibited bidirectional strengthening, whereas the pattern for semantically unrelated pairs was more consistent with a model in which forward (A-?) and backward (B-?) associations are strengthened independently from each other.

General discussion

When participants study a pair of words (A-B; e.g., BAKER-MOUNTAIN), they have to create an episodic association between them. There are two opposing views about the

directionality of these episodic associations: the independent-association hypothesis (IAH) posits that associations are unidirectional and separately modifiable ($A \rightarrow B$ and $A \leftarrow B$; Ebbinghaus, 1885/1913; Robinson, 1932), whereas the associative-symmetry hypothesis (ASH) proposes that only a single, bidirectional association exists between A and B ($A \leftrightarrow B$; Asch & Ebenholtz, 1962; Köhler, 1947). Our contribution is twofold. First, whereas the existing literature has focused on comparing the equality and correlation of forward (BAKER-?) and backward (MOUNTAIN-?) tests and favors ASH over IAH (Asch & Ebenholtz, 1962; Ekstrand, 1966; Horowitz, Norman, & Day, 1966; Kahana, 2002; Murdock, 1966), we provide here the first direct evidence for IAH, by showing that forward and backward associations are separately modifiable for semantically related pairs, but not for semantically unrelated pairs. Second, although most debates concerning ASH and IAH favor one hypothesis over the other, we show that episodic associations can be either independent or symmetric, depending on the semantic relatedness of the word pairs. Thus, whether or not episodic associations are symmetric depends on prior knowledge.

In two experiments, we presented participants with pairs of unrelated (e.g., BAKER-MOUNTAIN) and related (e.g., TEACHER-CLASSROOM) words in a study and testing paradigm in which repeated cued-recall test cycles used alternating associative directions: A-?, B-?, A-?, and so forth. We found that for unrelated words, performance improved between the tests, regardless of the change in direction. The restudy opportunity after every cued-recall test trial strengthened the bidirectional association, which is consistent with previous findings of associative symmetry in episodic memory for unrelated pairs (Caplan et al., 2006; Kahana, 2002; Madan et al., 2010). Similar results have previously been found for unrelated object-location pairs, in which an improvement in successive testing in alternating directions was taken as evidence for associative symmetry (Sommer et al., 2008). In contrast, semantically related pairs produced a strikingly different pattern—performance improved only between cycles that were testing memory in the same direction (i.e., RTs and accuracy improved across cycles with the same cue; e.g., Cycle 1 to 3 to 5 to 7), but performance was equivalent between tests that switched direction (e.g., performance on Cycle 1 = Cycle 2, Cycle 3 = Cycle 4, etc.). Although the stair-like pattern was less pronounced in Experiment 2, the key issue is whether performance improves when switching testing direction, which was the case only for semantically unrelated and not for semantically related pairs in both experiments. Thus, for semantically related pairs, when on Cycle 1 memory was tested with the A cue (followed by restudy), participants strengthened the representation of AB, but not that of BA, as shown through improvement on the next test of AB (Cycle 3), but not of BA (Cycle 2). The stair-like pattern for semantically related pairs is consistent with a model in

³ This analysis was not preregistered, since we did not expect performance to be at ceiling.

⁴ This analysis was also not preregistered.




Fig. 5 Distributions of differences in average response times (left) and cued-recall accuracy (right) between Test Cycles 2 and 1 in Experiment 2, for semantically related pairs. The red vertical lines show the average difference for semantically unrelated pairs. Since unrelated pairs were

randomly generated for each participant, there was generally only one observation per unrelated pair, which is why only the average difference is displayed for unrelated pairs (see the Materials section for more details on the randomization procedure)

which the pairs AB and BA are stored as two separate representations, which supports the IAH.

Interestingly, performance on the first two cycles was equivalent for semantically related pairs. This suggests that during the initial study phase, participants originally created representations for both AB and BA; otherwise, performance on Cycle 2 should have been at floor. This result is consistent with prior research that has repeatedly shown that backward tests are as accurate as forward tests when participants perform only one of those tests (Asch & Ebenholtz, 1962; Ekstrand, 1966; Horowitz, Norman, & Day, 1966; Murdock, 1966; for a review, see Kahana, 2002). As we discussed in the introduction, Wolford (1971) and Kahana (2002) argued that finding equal forward and backward recall is not inconsistent with the IAH. By extension, the equal recall on Cycles 1 and 2 is not inconsistent with the IAH, because the IAH requires that the two associations be independent, not that only one be created upon study. On the basis of these prior findings, we would expect to see the same pattern even if Cycle 1 had used the B word as a cue, and Cycle 2 used the A word as a cue, and so forth (as opposed to the present design, in which the first cycle tested pairs in the original study direction). We should also note that our main question was not whether forward and backward connections are both created, but whether they are separately modifiable. Combined with the stair-like main pattern, the fact that Cycle 2 performance was not at floor implies that whereas the initial study of semantically related pairs might create two associations in each direction of equal strength, these associations are independent, because restudy after testing only strengthens the association in the direction from the cue to the target, and not vice versa.

How generalizable are these findings? The first experiment had some issues that limited generalizability. First, we used double-function lists, in which each word appeared in more than one word pair (i.e., A–B, B–C, C–D), such that the shared cues could have created interference in retrieval from one test to the next. Second, the stimuli were only exemplars of person–

location pairs such as SCHOOL–TEACHER, which limited generalizability to other relations. In Experiment 2 we replicated the main findings while using only single-function lists (i.e., each word was presented in only one list), and with a wide variety of semantic relations (taxonomic, part of, composition, etc.; see Popov et al., 2017). Because of the single-function lists, accuracy was much higher in Experiment 2, and it reached ceiling levels for semantically related pairs after Test Cycle 3. Nevertheless, the lack of improvement between odd and even cycles for semantically related pairs was not due to ceiling effects—whereas accuracy for semantically related pairs did not change from Test 1 to Test 2, it did improve significantly in Test 3. In summary, the stair-like pattern that is consistent with associative asymmetry for semantically related pairs appears replicable and generalizable to a variety of relations.

Potential mechanisms

There are multiple reasons why semantically related items might show associative asymmetry. The two main mechanisms that we propose are unitization or encoding the preexisting semantic relation. According to the unitization account, people in our study might unitize semantically related pairs into single items—for example, SCHOOLTEACHER and TEACHERSCHOOL—which they might represent, access, and modify separately from one another. This process would be similar to the one proposed to explain asymmetry effects found for existing compound words (Caplan et al., 2014). This account bears some similarity to a proposal by Kahana (2002), who argued that asymmetry could arise if the presence of one word changes the interpretation of the other. As we noted in the introduction, the word HUNTER might have the same visual form in the pairs FOREST HUNTER and HUNTER FOREST, but it might represent different concepts in each pair—in the first pair it is an agent (a hunter who hunts in forests), whereas in the second pair, it is a modifier for the FOREST CONCEPT, noting that this is a

forest meant for hunting (Coutanche et al., [in press](#)). As a result, people might be retrieving different cognitive units with the same visual form on successive tests, explaining the lack of transfer between tests of different directions.

The second possibility is that in addition to the individual items, people store the specific semantic relation present in each pair (Popov et al., [2017](#)). The order of items matters for semantic relations—the complimentary relations *X works in Y* and *X is a workplace for Y* are potentially represented separately from one another. These separately represented relations can be utilized after system consolidation, in which newly learnt associations are encoded into the distributed cortical network with reduced involvement of the hippocampal system. It has been shown that rapid system consolidation can take place when the newly learnt information is consistent with schema (Tse et al., [2007](#)); the cue-recall test cycles further facilitate this process as it has been proposed that retrieval-based learning serves as a rapid consolidation event through reactivation of associative information (Antony, Ferreira, Norman, & Wimber, [2017](#)). These two representations are then strengthened separately and contribute independently to the retrieval of X when cued with Y, and vice versa. In contrast, unrelated items have to be encoded hippocampally, where the same pattern-completion process can occur regardless of which item is provided (see Caplan et al., [2014](#), for a similar proposal). Further research will be needed to discriminate between these possible explanations.

Finally, it is possible that participants use different strategies to process semantically related and unrelated pairs. For example, participants might be using rote rehearsal of the existing unidirectional semantic relationship to reinforce memory for semantically related pairs; in contrast, for semantically unrelated pairs participants might be actively constructing novel semantic associations between the words to aid subsequent retrieval. This processing difference in rote rehearsal versus active semantic elaboration is plausible, but in and of itself it does not explain why these strategies might lead to asymmetric versus symmetric associations. One could construct different theoretical models in which either strategy could lead to both types of associations (i.e., nothing dictates that participants would not rote rehearse both directions for semantically related pairs, or that they would not create separate semantic elaborations for each direction of semantically unrelated pairs). Thus, the type of processing itself cannot a priori explain whether an association would be created symmetrically or not. Nevertheless, it is possible to merge this account with either of the two mechanistic explanations proposed above—for instance, unitized items or the presence of a directional semantic relation could lead to rote rehearsal of the unit or the representation of the semantic relation, whereas items that cannot be unitized or for which no relational representation exists could promote active semantic elaboration instead, which would create symmetric associations by default.

Implications and open questions

The present results also have implications for retrieval practice in educational settings. It has often been observed that retrieval practice improves memory more than simply restudying information (the “testing effect”). Researchers have argued that using retrieval practice can improve memory in educational settings (McDaniel, Roediger, & McDermott, [2007](#); Roediger & Karpicke, [2006](#)). However, the majority of the data concerning the testing effect comes from studies using novel arbitrary associations (though see Carpenter, Pashler, & Cepeda, [2009](#)); in contrast, most real-world learning involves meaningfully related materials. The present results suggest that although testing might strengthen the association between a specific cue and the target, it will not necessarily improve overall access to the memory representation through multiple directions.

Future work could explore several questions. First, it would be valuable to contrast performance over successive tests in the alternating A-to-B and B-to-A testing used here with a condition in which each test cycle used the same cue (always A-? or B-?). Whereas here we contrasted related to unrelated pairs, in such a design, one could compare performance on the same set of items, depending on whether the tests changed direction or not: We would expect that unrelated pairs would improve gradually regardless of whether the tests changed direction; at the same time, performance for related pairs should improve gradually only in the condition in which all tests were in the same direction. If the forward and backward associations for related pairs are indeed independent, and their strength is not influenced by testing in the opposite direction, we would expect to see that performance on Test 2 in the nonalternating test condition should match performance on Test 3 in the alternating condition, performance on Test 3 in the nonalternating test condition should match performance on Test 5 in the alternating condition, and so forth. While providing such a comparison would add support to our conclusion, we believe that it is not strictly required. This is due to the dissociation between related pairs, which show the stair-like pattern, and the unrelated pairs, which show gradual improvement—these two conditions serve as controls for each other.

Second, in addition to testing the alternative mechanisms of unitization and processing of preexisting relations, future work could attempt to illuminate how these mechanisms interact with encoding and retrieval during testing. As we noted above, the fact that performance on Test 1 and Test 2 is equivalent for semantically related pairs suggests that, at least during the initial study sessions, participants create both associations/units (AB and BA), and only as a result of the subsequent testing do they begin to strengthen those independently. It is not currently clear whether retrieval of the stored unit or its subsequent reencoding during testing leads to the stair-like pattern we observed.

Our findings extend the results of Caplan et al. ([2014](#)), who found that associative symmetry is violated for pairs of items

that form compound words (e.g., ROSE–BUD). This raises the following question: Could other manipulations potentially lead to asymmetry in episodic associations? Models of episodic memory differ in the ways they represent episodic associations (as was reviewed by Kahana, 2002), and any violations of associative symmetry would present key data for testing such models. It is possible that some of the effects presented in this article could be due to the easier processing/better processing fluency of semantically related pairs, whereas the significantly more difficult processing of unrelated pairs caused participants to encode them and their associations in a qualitatively different manner. If this is the case, manipulations that increase the processing fluency of unrelated pairs (for a discussion of processing fluency, see Ahmad & Hockley, 2017) might also lead to violations in associative symmetry, despite the absence of a preexisting semantic relationship. In summary, the present demonstration of associative asymmetry opens up new avenues for future empirical research, model testing, and development.

Finally, although the present work has not addressed this, it is worth asking whether the episodic associations between items are represented as connections between them, or whether the associations are represented as items in their own right (essentially forming a triad in memory with the two items). Although Popov et al. (2017) has provided evidence that semantic relations are represented independently of the items that instantiate them, it is worth extending this question to episodic relationships. In some episodic memory models, such as SAC (Reder et al., 2000; Popov & Reder, 2018), experiencing concepts together creates a novel episode node in a network to which the two concept nodes are connected. These nodes have similar properties regardless of whether they are semantic or episodic in nature. In contrast, some distributed memory models represent associations as some transformation between item vectors, such as their outer-product or convolution (for a review, see Kahana, 2002).

Author note We thank John Paulus, Kimberly Hoover, and Chao Wu for assistance with the study. All of our stimuli, experimental software, data, and analysis code are publicly available on the Open Science Framework (<https://osf.io/72amw>), and Experiment 2 was preregistered (<https://osf.io/rdsw5>). We have also posted a preprint version of this manuscript at PsyArXiv, <https://psyarxiv.com/6pgjn/>.

References

- Ahmad, F. N., Fernandes, M., & Hockley, W. E. (2015). Improving associative memory in older adults with unitization. *Aging, Neuropsychology, and Cognition*, 22, 452–472.
- Ahmad, F. N., & Hockley, W. E. (2017). Distinguishing familiarity from fluency for the compound word pair effect in associative recognition. *Quarterly Journal of Experimental Psychology*, 70, 1768–1791
- Anderson, J. R., & Ross, B. H. (1980). Evidence against a semantic–episodic distinction. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 441–466. <https://doi.org/10.1037/0278-7393.6.5.441>
- Antony, J. W., Ferreira, C. S., Norman, K. A., & Wimber, M. (2017). Retrieval as a fast route to memory consolidation. *Trends in Cognitive Sciences*, 21, 573–576.
- Asch, S. E., & Ebenholtz, S. M. (1962). The principle of associative symmetry. *Proceedings of the American Philosophical Society*, 106, 135–163.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390–412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Badham, S. P., Estes, Z., & Maylor, E. A. (2012). Integrative and semantic relations equally alleviate age-related associative memory deficits. *Psychology and Aging*, 27, 141–152. <https://doi.org/10.1037/a0023924>
- Bartling, C. A., & Thompson, C. P. (1977). Encoding specificity: Retrieval asymmetry in the recognition failure paradigm. *Journal of Experimental Psychology: Human Learning and Memory*, 3, 690–700. <https://doi.org/10.1037/0278-7393.3.6.690>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67, 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bein, O., Reggev, N., & Maril, A. (2014). Prior knowledge influences on hippocampus and medial prefrontal cortex interactions in subsequent memory. *Neuropsychologia*, 64, 320–330.
- Caplan, J. B., Boulton, K. L., & Gagné, C. L. (2014). Associative asymmetry of compound words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40, 1163–1171. <https://doi.org/10.1037/a0036588>
- Caplan, J. B., Glaholt, M. G., & McIntosh, A. R. (2006). Linking associative and serial list memory: Pairs versus triples. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 1244–1265. <https://doi.org/10.1037/0278-7393.32.6.1244>
- Carpenter, S. K., Pashler, H., & Cepeda, N. J. (2009). Using tests to enhance 8th grade students’ retention of U.S. history facts. *Applied Cognitive Psychology*, 23, 760–771. <https://doi.org/10.1002/acp.1507>
- Coutanche, M. N., Solomon, S. H., & Thompson-Schill, S. L. (in press). Conceptual combination. In D. Poeppel, G. R. Mangun, & M. S. Gazzaniga (Eds.), *The cognitive neurosciences* (6th ed.). Cambridge, MA: MIT Press.
- Coutanche, M. N., & Thompson-Schill, S. L. (2015). Rapid consolidation of new knowledge in adulthood via fast mapping. *Trends in Cognitive Sciences*, 19, 486–488.
- De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2018). The “Small World of Words” English word association norms for over 12,000 cue words. *Behavior Research Methods*. Advance online publication. <https://doi.org/10.3758/s13428-018-1115-7>
- Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology* (H. A. Ruge & C. A. Bussenius, Trans.). New York, NY: Teachers College, Columbia University. (Original work published 1885)
- Ekstrand, B. R. (1966). Backward associations. *Psychological Bulletin*, 65, 50–64. <https://doi.org/10.1037/h0022645>
- Epstein, M. L., Phillips, W. D., & Johnson, S. J. (1975). Recall of related and unrelated word pairs as a function of processing level. *Journal of Experimental Psychology: Human Learning and Memory*, 1, 149–152. <https://doi.org/10.1037/0278-7393.1.2.149>
- Ford, J. H., Verfaellie, M., & Giovanello, K. S. (2010). Neural correlates of familiarity-based associative retrieval. *Neuropsychologia*, 48, 3019–3025.
- Greenberg, D. L., & Verfaellie, M. (2010). Interdependence of episodic and semantic memory: Evidence from neuropsychology. *Journal of*

- the International Neuropsychological Society, 16, 748–753. <https://doi.org/10.1017/S1355617710000676>
- Günther, F., Dudschig, C., & Kaup, B. (2015). LSAfun—An R package for computations based on Latent Semantic Analysis. *Behavior Research Methods*, 47, 930–944. <https://doi.org/10.3758/s13428-014-0529-0>
- Guttentag, R. (1995). Children's associative learning: Automatic and deliberate encoding of meaningful associations. *American Journal of Psychology*, 108, 99–114. <https://doi.org/10.2307/1423103>
- Haskins, A. L., Yonelinas, A. P., Quamme, J. R., & Ranganath, C. (2008). Perirhinal cortex supports encoding and familiarity-based recognition of novel associations. *Neuron*, 59, 554–560. <https://doi.org/10.1016/j.neuron.2008.07.035>
- Horowitz, L. M., Norman, S. A., & Day, R. S. (1966). Availability and associative symmetry. *Psychological Review*, 73, 1–15. <https://doi.org/10.1037/h0022661>
- Jones, L. L., Estes, Z., & Marsh, R. L. (2008). An asymmetric effect of relational integration on recognition memory. *Quarterly Journal of Experimental Psychology*, 61, 1169–1176.
- Kahana, M. J. (2002). Associative symmetry and memory theory. *Memory & Cognition*, 30, 823–840. <https://doi.org/10.3758/BF03195769>
- Kato, K., & Caplan, J. B. (2017). Order of items within associations. *Journal of Memory and Language*, 97, 81–102. <https://doi.org/10.1016/j.jml.2017.07.001>
- Köhler, W. (1947). *Gestalt psychology*. New York, NY: Liveright.
- Kriukova, O., Bridger, E., & Mecklinger, A. (2013). Semantic relations differentially impact associative recognition memory: Electrophysiological evidence. *Brain and Cognition*, 83, 93–103. <https://doi.org/10.1016/j.bandc.2013.07.006>
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–240. <https://doi.org/10.1037/0033-295X.104.2.211>
- Liu, X. L., Liang, P., Li, K., & Reder, L. M. (2014). Uncovering the neural mechanisms underlying learning from tests. *PLoS ONE*, 9, e92025. <https://doi.org/10.1371/journal.pone.0092025>
- Lockhart, R. S. (1969). Retrieval asymmetry in the recall of adjectives and nouns. *Journal of Experimental Psychology*, 79, 12–17. <https://doi.org/10.1037/h0026944>
- Madan, C. R., Glaholt, M. G., & Caplan, J. B. (2010). The influence of item properties on association-memory. *Journal of Memory and Language*, 63, 46–63. <https://doi.org/10.1016/j.jml.2010.03.001>
- McClelland, J. L. (2013). Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Journal of Experimental Psychology: General*, 142, 1190–1210. <https://doi.org/10.1037/a0033812>
- McCloskey, M., & Santee, J. L. (1981). Are semantic memory and episodic memory distinct systems? *Journal of Experimental Psychology: Human Learning and Memory*, 7, 66–71. <https://doi.org/10.1037/0278-7393.7.1.66>
- McDaniel, M. A., Roediger, H. L., III, & McDermott, K. B. (2007). Generalizing test-enhanced learning from the laboratory to the classroom. *Psychonomic Bulletin & Review*, 14, 200–206. <https://doi.org/10.3758/BF03194052>
- McKoon, G., Ratcliff, R., & Dell, G. S. (1985). The role of semantic information in episodic retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 742–751. <https://doi.org/10.1037/0278-7393.11.1.4.742>
- McKoon, G., Ratcliff, R., & Dell, G. S. (1986). A critical evaluation of the semantic-episodic distinction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 295–306. <https://doi.org/10.1037/0278-7393.12.2.295>
- Murdock, B. B., Jr. (1966). Forward and backward associations in paired associates. *Journal of Experimental Psychology*, 71, 732–737. <https://doi.org/10.1037/h0023113>
- Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, 36, 402–407. <https://doi.org/10.3758/BF03195588>
- Parks, C. M., & Yonelinas, A. P. (2015). The importance of unitization for familiarity-based learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41, 881–903. <https://doi.org/10.1037/xlm0000068>
- Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, 162, 8–13. <https://doi.org/10.1016/j.jneumeth.2006.11.017>
- Popov, V., & Hristova, P. (2015). Unintentional and efficient relational priming. *Memory & Cognition*, 43, 866–878. <https://doi.org/10.3758/s13421-015-0514-6>
- Popov, V., Hristova, P., & Anders, R. (2017). The relational luring effect: Retrieval of relational information during associative recognition. *Journal of Experimental Psychology: General*, 146, 722–745. <https://doi.org/10.1037/xge0000305>
- Popov, V., & Reder, L. (2018). *Frequency effects on memory: A resource-limited theory*. OSF preprint. <https://doi.org/10.31219/osf.io/dsx6y>
- R Core Team. (2014). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from www.R-project.org
- Reder, L. M., Nhouyvanisvong, A., Schunn, C. D., Ayers, M. S., Angstadt, P., & Hiraki, K. (2000). A mechanistic account of the mirror effect for word frequency: A computational model of remember-know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 294–320. <https://doi.org/10.1037/0278-7393.26.2.294>
- Rehani, M., & Caplan, J. B. (2011). Interference and the representation of order within associations. *Quarterly Journal of Experimental Psychology*, 64, 1409–1429. <https://doi.org/10.1080/17470218.2010.549945>
- Rhodes, S. M., & Donaldson, D. I. (2007). Electrophysiological evidence for the influence of unitization on the processes engaged during episodic retrieval: Enhancing familiarity based remembering. *Neuropsychologia*, 45, 412–424.
- Robinson, E. S. (1932). *Association theory to-day: An essay in systematic psychology*. New York, NY: Century Co.
- Roediger, H. L., III, & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17, 249–255. <https://doi.org/10.1111/j.1467-9280.2006.01693.x>
- Roediger, H. L., III, & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 803–814. <https://doi.org/10.1037/0278-7393.21.4.803>
- Rowland, C. A. (2014). The effect of testing versus restudy on retention: A meta-analytic review of the testing effect. *Psychological Bulletin*, 140, 1432–1463. <https://doi.org/10.1037/a0037559>
- Silberman, Y., Bentin, S., & Miikkulainen, R. (2007). Semantic boost on episodic associations: An empirically-based computational model. *Cognitive Science*, 31, 645–671. <https://doi.org/10.1080/15326900701399921>
- Slamecka, N. J. (1976). An analysis of double-function lists. *Memory & Cognition*, 4, 581–585. <https://doi.org/10.3758/BF03213221>
- Sommer, T., Rose, M., & Büchel, C. (2007). Associative symmetry versus independent associations in the memory for object-location associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 90–106. <https://doi.org/10.1037/0278-7393.33.1.90>
- Sommer, T., Schoell, E., & Büchel, C. (2008). Associative symmetry of the memory for object–location associations as revealed by the testing effect. *Acta Psychologica*, 128, 238–248. <https://doi.org/10.1016/j.actpsy.2008.01.003>

- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. M. (2007). Schemas and memory consolidation. *Science*, *316*, 76–82. <https://doi.org/10.1126/science.1135935>
- van Kesteren, M. T. R., Ruiter, D. J., Fernández, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends in Neurosciences*, *35*, 211–219. <https://doi.org/10.1016/j.tins.2012.02.001>
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, *14*, 779–804. <https://doi.org/10.3758/BF03194105>
- Wolford, G. (1971). Function of distinct associations for paired-associate performance. *Psychological Review*, *78*, 303–313. <https://doi.org/10.1037/h0031032>
- Zhang, Q., Popov, V., Koch, G. E., Calloway, R. C., & Coutanche, M. N. (2018). Fast memory integration facilitated by schema consistency. In C. Kalish, M. Rau, J. Zhu, & T. T. Rogers (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 2777–2782). Austin, TX: Cognitive Science Society. <https://doi.org/10.1101/253393>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.