A Specificity Principle of Memory: Evidence from Aging and Associative Memory

Nathaniel R. Greene & Moshe Naveh-Benjamin

University of Missouri

Author Note

Nathaniel R. Greene and Moshe Naveh-Benjamin, Department of Psychological Sciences, University of Missouri.

Correspondence concerning this article should be addressed to Nathaniel R. Greene, Department of Psychological Sciences, University of Missouri, Columbia, MO 65211; E-mail: ngwfm@mail.missouri.edu, or to Moshe Naveh-Benjamin, Department of Psychological Sciences, University of Missouri, 106 McAlester Hall; E-mail: navehbenjaminm@missouri.edu.

Data and material for this study are openly available at the Open Science Framework, <https://osf.io/xk78c/>.

N. R. Greene was supported by a Department of Psychological Sciences Graduate Research Grant.

This is the accepted, pre-production version of an article in-press with *Psychological Science*, not the official copy of record. Please contact one of the authors before citing this version.

Abstract

The ability to remember associations among components of an event, which is central to episodic memory, declines with normal aging. In accord with the specificity principle of memory (Surprenant & Neath, 2009), these declines may occur because associative memory requires retrieval of specific information. Guided by this principle, we endeavored to determine whether ubiquitous age-related deficits in associative memory (e.g., Naveh-Benjamin, 2000) are restricted to specific representations or extend to the gist of associations. Young and old adults (30 each in Experiment 1, 40 each in Experiment 2) studied face-scene pairs and were administered associative recognition tests following variable delays. Whereas both young and older adults could retrieve the gist of associations, older adults were impaired in their ability to retrieve more specific representations. Our results also show that associations can be retrieved from multiple levels of specificity, suggesting that episodic memory might be accessed on a continuum of specificity.

Keywords: *memory, aging, associative recognition, conjoint recognition, principles of memory, multinomial processing tree models*

A Specificity Principle of Memory: Evidence from Aging and Associative Memory

Scientific principles describe invariants of natural phenomena, but efforts to identify principles of memory have been scarce and often unsuccessful (Roediger, 2008). However, Surprenant and Neath (2009) provided compelling evidence that several principles may exist across all memory systems. Since their book was published, there have been relatively few attempts to test their proposed principles and fewer attempts to assess if these principles can explain changes in memory performance that occur with normal aging. Here, we report on their suggested specificity principle, according to which vulnerability to forgetting is highest for tasks requiring individuals to retrieve specific information.

The specificity principle may provide insight into causes of pervasive age-related deficits in episodic memory, particularly for the associations among components of an episode (e.g., Naveh-Benjamin, 2000; Old & Naveh-Benjamin, 2008). Episodic memories are bounded representations of contextual, temporal, and spatial components (Tulving, 1983). According to the associative deficit hypothesis (Naveh-Benjamin, 2000), older adults can remember individual components of an episode relatively well, compared to young adults, but often fail to encode or retrieve the associations among these components. These failures contribute to memory errors like misremembering where a person was previously encountered (Chen & Naveh-Benjamin, 2012) and can explain why older adults often forget the source of information (Boywitt, Kuhlmann, & Meiser, 2012). We propose that at least some of these deficits arise because associative memory tasks require individuals to retrieve specific information, a process impaired by normal aging.

Traditionally, associative memory tasks require individuals to distinguish previously studied associations (A-B and C-D) from recombined associations (A-C or B-D). To be

successful at this task, one must retrieve some specific information about which items were associated at encoding. Because older adults are prone to associative memory errors, an important unanswered question remains: How *much* specific information about associations can older adults retrieve? That is, do they fail to retrieve specific bounded representations, or do these deficits also extend to retrieval of the gist of associations?

The distinction between specific and gist memory is central to Fuzzy-Trace Theory (FTT; Reyna & Brainerd, 1995), which postulates that the surface-level details and the meaning of an episode are processed in parallel in a verbatim and gist trace, respectively. FTT is commonly used to explain false memory phenomena, like those observed in the Deese-Roediger-McDermott paradigm (Roediger & McDermott, 1995), where individuals study a list of items (e.g., “dream”, “rest”, and “bed”) that are closely related to an unpresented lure (“sleep”). False recognition of the lure is well-documented in older adults, which is taken as evidence that they use a gist-based processing strategy (for a recent review, see Devitt & Schacter, 2016).

Falsely recognizing a lure in this paradigm suggests that older adults remember the *semantic gist* of the studied material (e.g., “dream,” “rest,” and “bed” are related to “sleep”), but it does not necessarily indicate that they retrieve a fuzzier episodic representation of a particular event. A more direct way to measure this is to assess specific versus gist memory for the associations among components of an episode that lie at the core of episodic memory.

There have been a few attempts to assess if older adults can remember the gist of associations, such as studies showing that they can retrieve the gist of item-price pairs (Castel, 2005; Castel, McGillivray, & Worden, 2013). However, it is possible that older adults could remember the gist of what price was associated with a given item in those studies because they could rely on pre-existing knowledge (e.g., “a gallon of milk is usually \$2 to \$3”). This may not

be possible for other associations, such as face-scene pairs. Although it is conceivable that older adults could use a pre-existing schema to remember the gist of where they saw a face (e.g., “this old man was in some park”), it is also possible that alternative schemas might interfere with this ability (e.g., “there may have been a young woman in this park”), which might result in failures to retrieve the gist of an association.

Further, it has not been determined whether young or older adults can retrieve associations with varying amounts of specificity. For example, memory of an old man in a park might be retrieved with high specificity (“he was in *this* park”), with less specificity (“he was in *some* park”), or with only a fuzzy representation reinstated (“he was somewhere in nature”).

Although FTT alludes to the possibility of hierarchies of gist (Brainerd & Reyna, 2015; Reyna & Brainerd, 1991), these hierarchies mostly pertain to semantic representations. In contrast, Craik (2002) proposed that memory might exist on a continuum with no categorical break between episodic and semantic representations. By this account, there could be multiple levels of specificity for an individual episode. We provide the first test of this proposal in the context of associative memory.

The Present Study

In two experiments, we tested the hypothesis that age-related associative memory deficits result, in part, from inabilities in retrieving specific representations, but that older adults could retrieve these associations at less specific levels of representation. We used face-scene pairs, which simulate remembering where someone was encountered (Grupposo, Lindsay, & Masson, 2007), as older adults have demonstrated an associative deficit for face-scene pairs in both short- and long-term memory (Chen & Naveh-Benjamin, 2012). To measure specific and gist memory, we used the Simplified Conjoint Recognition (SCR) paradigm and its multinomial processing

tree model (Stahl & Klauer, 2008). Although the paradigm is based on FTT, in Experiment 2, we modified it to account for the possibility of different levels of specificity. Finally, principles of memory are proposed to operate on all time-scales (Surprenant & Neath, 2009), so we measured specific and gist memory for associations following various retention intervals between encoding and retrieval.

Experiment 1

Method

Participants. Thirty young adults, 60% female, aged 18 to 23 ($M = 19.00$, $SD = 1.14$), participated in exchange for psychology research credits. Thirty old adult participants, 70% female, aged 62 to 80 ($M = 72.83$, $SD = 4.39$) were recruited from the surrounding community and compensated \$15 for their time and travel expenses. All older adults reported to be in good health, had normal or corrected-to-normal hearing and vision, and scored 27 or higher on the Mini-Mental State Examination. As analyses were based on Bayesian posterior estimation, a formal power analysis, which is non-Bayesian by nature, was not conducted, given that evidence for a null effect is equally informative in the Bayesian framework (e.g., Kruschke, 2011, 2018). Instead, we elected to use a sample size of 30 participants per age group, as this compares favorably with the sample sizes in the original SCR paradigm (Stahl & Klauer, 2008). Upon reaching this sample size, we used a Bayesian prior sensitivity analysis, specifying our multinomial processing tree models with priors that varied across three levels of informativeness. This type of analysis tests whether the results are robust to choice of prior. In our case, they were (see <https://osf.io/fpcmu/>). It is important to note that optional stopping (e.g., terminating data collection at our elected sample size) does not bias the obtained results in the Bayesian framework (Rouder, 2014).

Stimuli. We paired 168 faces with 168 scenes. Faces were drawn from the FACES database (Ebner, Riediger, & Lindenberger, 2010), with 42 faces from each of four age and gender normed categories (i.e., young female, young male, old female, old male). Scenes were drawn from a categorized scene pool (Konkle, Brady, Alvarez, & Olivia, 2010), with two exemplars from each of 84 categories (e.g., two parks, two malls).

We divided 144 face-scene pairs evenly across 12 experimental blocks, and the remaining 24 pairs were presented in two practice blocks. Each block featured six face-scene pair “categories,” in which both exemplars of a given scene category were paired with faces from the same age and gender category. For example, two parks were paired with two old male faces in one study block. This manipulation was designed to enhance schematic support for a given association (e.g., “parks are paired with old men”) to potentially reduce the influence of other pre-existing schemas (e.g., “young men, young women, or old women could also be in parks”). Participants thus needed to maintain 12 unique associations (e.g., Old Man 1 – Park 1, Old Man 2 – Park 2) but only six gist-based categorized associations (e.g., Old Man – Park).

Pairs were automated using E-Prime 2.0 software and presented on a 20-inch ASUS flat-screen LED monitor with a resolution of 1920 x 1080. Faces appeared to the left of the center of the screen and scenes to the right. We also constructed a list of 136 populous United States cities and their corresponding states for a verbal semantic memory filler task on blocks featuring a longer delay between the study and test phases.

Procedure. All participants provided informed consent prior to participation. Participants were tested individually with the experimenter present. The experimental session lasted about an hour and was divided into two practice and 12 experimental blocks. Each block consisted of a study phase and a test phase (see Figure 1a). During the study phase, face-scene pairs were

presented sequentially onscreen at a rate of one pair every four seconds, with a 0.5 second inter-pair interval. Participants were instructed to study the pairs for a forthcoming memory test. They were informed that no two pairs would be identical but that some pairs might appear to be quite similar. Twelve face-scene pairs from the set of all possible pairs were pseudo-randomly intermixed during a given study block, with the only constraint being that both pairs from a given categorized pairing had to appear in the same block. In six of the blocks, the test phase followed approximately five to eight seconds after the last pair disappeared from the screen (short delay), while in the remaining six blocks, the test phase occurred approximately two and a half minutes later following a verbal semantic memory task (long delay).

In the short delay blocks, the word “Wait” appeared, centered on screen, for five seconds between the study and test phases, and was followed by a test prompt informing participants the test phase was about to begin. Participants pressed the spacebar to initiate the test phase, or it began automatically after three seconds. In the long delay blocks, an instruction prompt for the filler activity appeared after the end of the study phase. Participants initiated the task by pressing the spacebar, or it began automatically after five seconds. For the semantic memory task, the names of 21 United States cities appeared on screen sequentially one at a time for five seconds, and participants were instructed to say aloud the state in which each city was located. The correct answer was shown for an additional two seconds. At the end of the task, a test prompt appeared as in the short delay blocks and was terminated with a spacebar press or after three seconds. Short and long delay blocks were randomly intermixed for each participant to ensure that participants would not know which delay to expect on a given block.

a

Experiment 1

b

Experiment 2

Fig. 1. Illustration of the procedure for Experiment 1 (a) and Experiment 2 (b).

During the test phase of each block, 12 pairs were presented at random, one at a time and evenly distributed into Intact, Related, and Unrelated probes. Intact test pairs featured a face-scene pair that had previously been presented during the study phase, such as the young woman paired with the desert in Figure 1a. Related test pairs were recombined face-scene pairs in which the face from one pair was recombined with a scene within the same category. For example, the Related probe depicted in the test phase of Figure 1a features an old man recombined with another kitchen scene from the study phase. Unrelated test pairs were recombined face-scene pairs in which the face from one pair was recombined with a scene from a different category, such as the young woman paired with the kitchen scene in the test phase of Figure 1a. There were 48 of each test probe throughout the experiment, evenly divided between the short and long delay blocks. Participants indicated whether a test pair was “intact,” “related,” or “unrelated” by pressing the “a,” “g,” or “l” key, respectively, which resulted in the next pair appearing. They were untimed in their responses. A three second inter-block interval separated the blocks. All procedures were approved by the University of Missouri’s Institutional Review Board.

Results

All data and analysis scripts are available at <https://osf.io/xk78c/>. Analyses were based on Bayesian posterior estimation, which confers many advantages over null hypothesis significance testing, such as the ability to quantify evidence for a null effect. Participants’ responses to each probe are depicted in Figure 2. As depicted, there was evident heterogeneity across participants, which was confirmed with chi-square tests for participant heterogeneity (Smith & Batchelder, 2008) fit separately to young and old adult data, all $\chi^2(145) \geq 627.13, p < .001$. Therefore, we used hierarchical models for both the logistic regression and multinomial processing tree (MPT) analyses reported below.

Experiment 1 Responses

a

Intact Probes

b

Related Probes

[Continued on next page]

Fig. 2. Proportion of responses to Intact probes (a), Related probes (b), and Unrelated probes (c) for Experiment 1. Lines at the top of the solid bars denote the group mean. Box around the mean indicates the 95% confidence interval of the mean. Jittered points denote individual participants' data.

Logistic regression results. For each probe, we measured age (coded as 0 = young, 1 = old) and delay (0 = short, 1 = long) differences in response accuracy (0 = incorrect, 1 = correct) with hierarchical Bayesian binomial logistic regression models, a more appropriate method of analyzing proportion correct data than ANOVA on aggregated proportions (Dixon, 2008). We also measured age and delay differences in response tendency to each probe (i.e., whether the participant responded “intact,” “related,” or “unrelated”), with categorical logistic regression models. Models were implemented in the *brms* package for R (Bürkner, 2017; R Core Team, 2018). Cauchy priors with location parameters of 0 and scale parameters of 2.5 (for the population-level slopes) and 10 (for the intercepts) were specified (see Gelman, Jakulin, Pittau, & Su, 2008), but we retained the default *t*-distribution priors on the random effects. Sampling routine information is available in the supplemental material (see section “Hierarchical Bayesian Categorical Logistic Regression Model Diagnostics”).

All parameters in logistic regression are on the log-odds scale, so we took the exponent of each estimate to return more easily interpretable odds-ratios (OR) and 95% Bayesian credible intervals (CI). The Bayesian CI conveys the most probable range of values of the true estimated OR, such that values within the interval have higher credibility than values outside the interval (Kruschke, 2011, 2018). Following the approach to parameter interpretation advocated by Kruschke (2011, 2018), we specified a region of practical equivalence (ROPE) to define a decision threshold to accept or reject a null value. It is important to note that Bayesian inference is concerned with describing our uncertainty in a parameter, rather than hard “black-and-white” decisions to accept or reject values. Nevertheless, we defined a null value (e.g., a null effect of age) as an OR of 1, with a ROPE ranging from 0.94 to 1.06. If the resulting 95% CI excludes the

ROPE, we reject the null. If the 95% CI is entirely contained within the ROPE, we accept the null. Otherwise, we remain agnostic.

First, we assessed whether there were any age differences in accuracy on the task. For Intact probes (Figure 2a), the resulting OR comparing old to young adults was 0.63 (95% CI: 0.42, 0.94), indicating that older adults were less accurate. This was likely attributable to older adults' greater tendency (relative to young adults) to respond "related" to Intact probes, OR = 1.55 (95% CI: 1.07, 2.27). The evidence was inconclusive as to whether there were age differences in the tendency to respond "unrelated," OR = 1.57 (95% CI: 0.84, 3.06), as there were relatively few "unrelated" responses in both age groups. Older adults also responded less accurately to Related probes (Figure 2b), OR = 0.36 (95% CI: 0.25, 0.52). Relative to young adults, older adults were more likely to respond that Related probes were "intact," OR = 2.86 (95% CI: 1.99, 4.06), or "unrelated," OR = 2.59 (95% CI: 1.40, 4.57). Finally, for Unrelated probes (Figure 2c), the OR = 0.61 (95% CI: 0.27, 1.30) partially overlapped with the ROPE and partially excluded it, such that we remain agnostic.

Next, regarding delay effects, the evidence for Intact probes was inconclusive, OR = 0.90 (95% CI: 0.76, 1.06), but note that the modal value lies quite near the ROPE, such that there was most probably very little, if any, change in accuracy across delay. However, accuracy for Related probes declined modestly with delay, OR = 0.77 (95% CI: 0.66, 0.90), with participants being more likely in the long delay blocks to endorse Related probes as "intact," OR = 1.39 (95% CI: 1.17, 1.63). In contrast, for Unrelated probes, accuracy slightly improved from short to long delay blocks, OR = 1.36 (95% CI: 1.11, 1.68), driven by a decreased tendency in the long delay blocks to endorse Unrelated probes as "related," OR = 0.74 (95% CI: 0.58, 0.95).

Finally, we tested if there were any Age by Delay interactions. In all cases, the 95% CI partially overlapped with the ROPE, but tended to exclude the ROPE in the tails of the posterior distributions. Therefore, we remain undecided, but a visual inspection of Figure 2 reveals very similar patterns in both the short and long delays.

Multinomial processing tree (MPT) results. The presence of notable age differences on Intact and Related probes, but not Unrelated probes, suggests that older adults might be deficient relative to younger adults at retrieving specific, but not gist, representations of associations. To test this assumption, we used the MPT model from the SCR paradigm (Stahl & Klauer, 2008), depicted in Figure 3. MPT models attempt to explain how participants arrive at their responses (boxes on the right) to a given memory probe (boxes on the left) by way of different cognitive processes (circles). The probability of making response j to probe i is expressed as the sum of the probabilities of the branches terminating in that response.

In Figure 3, the two V parameters correspond to the probability that participants retrieve the specific (verbatim) representation of an association given either an Intact probe (V_i) or a Related probe (V_r). The two G parameters correspond to the conditional probabilities that participants retrieve the gist of an association, for Intact (G_i) or Related (G_r) probes, given that they have not retrieved its more specific representation. If participants retrieve the gist, they then guess whether the probe is “intact” (with probability a) or “related” (with complementary probability $1-a$). If a probe elicits no verbatim or gist information for a participant, then the participant can still guess that the probe is “intact” or “related” with probability b . Otherwise, with probability $1-b$, the participant responds that the probe is “unrelated.” Note that the Unrelated probes contain no verbatim or gist parameters because these probes contain no verbatim or gist representation of originally-studied associations, and thus responses to Unrelated

probes are based entirely on parameter b . Conditional on entering cognitive state b , the participant guesses whether the probe is “intact” (with probability a_b) or “related” (with probability $1-a_b$). The original SCR paradigm (Stahl & Klauer, 2008) contained a single a parameter for responding “old” whether or not the gist was retrieved. However, this model fit our data poorly (see supplement for details), so we included the additional a_b parameter, which assumes that participants’ bias to respond “intact” can change depending on whether or not the gist is retrieved. This modified model had more parameters than degrees of freedom in the data, so we imposed an equality restriction on the gist parameters, as the scene paired with a face for both Intact and Related probes was drawn from the same category. No further parameter restrictions were warranted as each parameter in the model measures a cognitive process of interest, and the model is identified.

We fit the model separately to both age groups, with hierarchical latent-trait priors (Klauer, 2010) specified using the TreeBUGS package for R (Heck, Arnold, & Arnold, 2018; R Core Team, 2018). In both groups, the 95% CIs for the difference in each parameter between the short and long delay blocks overlapped with 0, so we collapsed delay within each group. This approach follows the interpretation of testing for group (or other factor) differences in MPT parameters advocated for by Smith and Batchelder (2010), who argued that if the difference was non-distinguishable from 0 (i.e., its credible interval overlapped with 0), then the parameter does not differ meaningfully across groups. Of course, one could argue that a ROPE may offer a more detailed test (e.g., Kruschke, 2018), but parameters in MPTs are correlated and inter-dependent, such that a ROPE is likely invalid in this context.

[Caption on next page]

Fig. 3. Multinomial processing tree model from the Simplified Conjoint Recognition paradigm (Stahl & Klauer, 2008). a = guessing “intact” upon retrieving the gist. a_b = guessing “intact” upon entering state b . b = guessing “intact or related.” G_i = probability the gist trace is retrieved given an Intact probe. G_r = probability the gist trace is retrieved given a Related probe. V_i = probability the verbatim trace is retrieved given an Intact probe. V_r = probability the verbatim trace is retrieved given a Related probe.

Table 1 describes the parameter estimates for each age group. For each parameter, we subtracted the posterior samples of the old from the young adults, to compute a credibility interval of the difference between groups (e.g., Smith & Batchelder, 2010). These differences are presented in Figure 4. Parameters for which the 95% CI of the difference estimate overlaps with 0 do not meaningfully differ between groups. As shown, young adults had higher values of both verbatim memory parameters (V_i and V_r) but did not differ from the old adults on gist memory parameters (G_i and G_r) or in response bias parameters (a , a_b , or b).

Because the gist parameters are *conditionally* reached upon failing to access the verbatim trace, the equivalent gist parameters indicate that both young and old adults are equally likely to retrieve the gist of an association when they fail to access the verbatim trace. However, to derive a composite score of how often participants retrieve some degree of representation for an original association, we must combine information about how often participants access *either* the verbatim or the gist traces. This can be achieved with the following equations: for Intact probes, $V_i + (1-V_i)*G_i$; for Related probes, $V_r + (1-V_r)*G_r$. Plugging in parameter estimates from Table 1, we derive composite retrieval scores for the Intact probes of .95 (95% CI: .89, .98) for the young adults and .91 (95% CI: .85, .95) for the old adults. For the Related probes, composite retrieval scores are .94 (95% CI: .89, .97) for the young adults and .90 (95% CI: .85, .94) for the old adults. That is, older adults can retrieve at least some degree of representation of associations as well as the young adults, but they are impaired in their ability to retrieve the most specific levels of representation.

Table 1. Multinomial processing tree parameter estimates.

Parameter	Experiment 1		Experiment 2	
	Younger Adults	Older Adults	Younger Adults	Older Adults
V_i	.39 [.18, .56]	.10 [.01, .24]	.59 [.49, .68]	.27 [.14, .39]
V_r	.36 [.21, .48]	.04 [.00, .12]	.27 [.14, .38]	.06 [.00, .13]
$G_i = G_r$	.91 [.86, .95]	.90 [.85, .93]	.63 [.56, .70]	.59 [.52, .64]
a	.63 [.53, .72]	.67 [.59, .75]	.64 [.51, .78]	.62 [.53, .72]
ab	.34 [.20, .51]	.35 [.25, .45]	.21 [.06, .43]	.29 [.25, .45]
b	.10 [.05, .17]	.15 [.10, .22]	.12 [.06, .17]	.28 [.21, .35]
F	----	----	.08 [.03, .13]	.05 [.01, .13]

Note: Estimates depict group level means and 95% Bayesian credible interval (in brackets). a = guessing “intact” upon retrieving the gist. a_b = guessing “intact” upon entering state b . b = guessing “intact or related.” F = a fuzzy level of representation is retrieved. G_i = gist trace is retrieved given an Intact probe. G_r = gist trace is retrieved given a Related probe. V_i = verbatim trace is retrieved given an Intact probe. V_r = verbatim trace is retrieved given a Related probe.

Fig. 4. Forest plots depicting the difference scores (Young minus Old) for each parameter in Experiment 1 (a) and Experiment 2 (b). Points denote the posterior mean difference, lines the 95% Bayesian credible interval of the difference score. Dashed line at 0 denotes a point-null. *a* = guessing “intact” when gist is retrieved; *ab* = guessing “intact” when gist not retrieved, *b* = guessing “intact” or “related,” *Gi* = gist is retrieved given an Intact probe, *Gr* = gist is retrieved given a Related probe, *F* = a fuzzy level of representation is retrieved, *Vi* = verbatim trace is retrieved given an Intact probe, *Vr* = verbatim trace is retrieved given a Related probe.

Signal detection theory (SDT) analyses. In addition to the MPT results described above, we also modeled our data using the equal variance SDT model, given an ongoing debate about whether recognition proceeds on a discrete or continuous basis (Batchelder & Alexander, 2013; Dubé & Rotello, 2012; Pazzaglia, Dubé, & Rotello, 2013). This debate has been somewhat extended to associative recognition (see Rotello, 2017) and to conjoint recognition tasks (Brainerd, Gomes, & Moran, 2014). The SDT results are presented in detail in the supplement and generally converge with our MPT results in both experiments (see Table S1). One notable exception is that the SDT results suggested that older adults in Experiment 1 had a more liberal bias (i.e., more likely to respond “intact”) when comparing Intact to Related probes, whereas the MPTs showed no age-related differences in response bias parameters. This is likely attributable to the fact that, in the MPT model, there are several processes that contribute to responding “intact,” including retrieving verbatim memory, guessing “intact” when only the gist is retrieved, and guessing “intact” when neither the verbatim nor the gist is retrieved. In contrast, in our SDT models, the only parameters are memory strength and response bias.

Discussion

These results support our hypothesis that age-related deficits in associative memory are restricted to the retrieval of specific, but not gist, representations. Older adults made more errors than young adults on Intact and Related, but not Unrelated probes, and the fit of an MPT model demonstrated that older adults had lower estimates of verbatim, but not gist, memory.

There was no effect of delay on verbatim or gist memory for associations. It is possible that the retention intervals used in the long delay blocks were not long enough to show any continued loss of verbatim memory or potential loss of gist memory, so we extended the retention interval in Experiment 2.

We categorized face-scene pairs to strengthen schematic support for a given association by offsetting alternative schemas, which might help with retrieving the gist. However, this created an imbalance in the number of specific associations participants needed to maintain per block (12; e.g., Old Man 1 – Park 1 and Old Man 2 – Park 2) compared to the number of gist associations to be maintained (6; e.g., Old Man - Park). Therefore, in Experiment 2, we remove categorized pairings. We also introduce a novel manipulation to measure different levels of specificity for associations.

Experiment 2

Method

Participants. We recruited 40 young adult participants, 72.5% female, aged 18 to 24 ($M = 19.68$, $SD = 1.86$), and 40 old adult participants, 75% female, aged 64 to 80 ($M = 73.05$, $SD = 3.99$). Young adults were compensated with research credits for psychology courses, and old adults were compensated with \$15 for time and travel expenses. Data from an additional 3 young and 4 old adults are not reported in these analyses as they performed well below chance level (i.e., below 33% accuracy across the different probe conditions). Our sample size was based on the same approach used in Experiment 1, but because the MPT parameters were estimated from fewer trials per participant, we increased our sample size to 40 participants to more precisely estimate the parameters. The parameter estimates were robust across different Bayesian prior specifications (see <https://osf.io/fpcmu/>).

Stimuli and Procedure. Stimuli were identical to those used in Experiment 1, but we created only 84 face-scene pairs, with 72 pairs presented during the experimental blocks and 12 in a practice block. Although the procedure was mostly similar to Experiment 1 (see Figure 1b), there were a few noteworthy changes. First, we reduced the total number of blocks from 12 to six

to keep the experimental session to under an hour. Three short and three long delay blocks were randomly intermixed for each participant, and the verbal semantic memory filler task was extended to occupy 301 seconds (approximately 5 minutes) between the study and test phases in the long delay blocks. This reduced the number of test probes to 24 of each type (i.e., Intact, Related, and Unrelated), evenly divided between short and long delay.

Second, we equated the number of specific and gist associations to be maintained in a single block by removing categorized pairings. Although two exemplars of a given scene type (e.g., two parks; see Figure 1b) both appeared in the same block, we ensured that each exemplar was paired with a different age or gender categorized face (e.g., Old Woman 1 – Park 1, Old Man 1 – Park 2). However, Related test pairs were still constructed by recombining a face with a similar scene (e.g., Old Woman 1 – Park 2; see test phase of Figure 1b).

Third, we organized the scenes in each block such that half of the scenes belonged to a broader nature category (e.g., two parks, two forests, two streams) and half to a broader indoor category (e.g., two corridors, two lobbies, two dens). This manipulation was designed to test the possibility for retrieval of different levels of specificity for associations, that is whether an individual can retrieve an original association (e.g., “this old man was with this park”), a less specific representation that still retains information about the type of scene paired with a given face (e.g., “this old man was with some park”), or a fuzzier representation retaining information about the broader scene category paired with a given face (e.g., “this old man was outside somewhere”). To test this, we constructed the Unrelated test pairs such that half of them featured a scene switch from within the same broader category (e.g., from Old Man 1 - Park 2 to Old Man 1 - Stream 1, both nature scenes; Unrelated-Within category) and half featured a scene switch

from outside the broader category (e.g., from Young Woman 1 - Stream 1 to Young Woman 1 - Lobby 2; Unrelated-Opposite category). The procedure was otherwise identical to Experiment 1.

Results

Logistic regression results. The proportion of responses to each probe is depicted in Figure 5. There was substantial heterogeneity in responses in both age groups, all $\chi^2(312) \geq 636.21$, $p < .001$. The age effects were mostly consistent with Experiment 1. Specifically, older adults were less accurate than young adults on Intact probes, OR = 0.47 (95% CI: 0.32, 0.69). However, whereas in Experiment 1 this was driven by older adults' tendency to endorse Intact probes as "related," in Experiment 2, they were not only more likely than young adults to respond "related," OR = 2.48 (95% CI: 1.57, 3.93), but they were also more likely to respond "unrelated" to Intact probes, OR = 1.93 (95% CI: 1.25, 3.00). As in Experiment 1, older adults were less accurate than young adults on Related probes, OR = 0.48 (95% CI: 0.36, 0.64), and were more likely than young adults to endorse Related probes as either "intact," OR = 2.20 (95% CI: 1.62, 2.97), or "unrelated," OR = 1.84 (95% CI: 1.36, 2.44). Unlike in Experiment 1, there was clear evidence in Experiment 2 that older adults were less accurate than young adults on Unrelated probes, OR = 0.45 (95% CI: 0.28, 0.73), and were more likely than young adults to endorse Unrelated probes as "intact," OR = 2.83 (95% CI: 1.58, 5.10), or "related," OR = 1.93 (95% CI: 1.12, 3.42).

For the Unrelated probes, we also tested if performance differed between the two types of Unrelated probes (coded as 0 = Unrelated-Within, 1 = Unrelated-Opposite). The age effect held for both types. However, across both age groups, accuracy was lower for the Unrelated-Within probes, OR = 0.61 (95% CI: 0.49, 0.76). This could reflect that some participants retrieved a fuzzier level of representation for these probes, which maintain some gist-like similarity to the

original associations. We test this possibility with the MPT analyses described in the next section. Finally, for all probes, we remained agnostic about any effects of delay or Age by Delay interactions, as in all cases, the 95% CI for the odds-ratio mostly overlapped with the ROPE, but the tails of the posterior distributions tended to marginally exclude the ROPE. A visual inspection of Figure 5 reveals very little, if any, changes in performance across delay.

Experiment 2 Responses

a

b

[Continued on next page]

c

d

Fig. 5. Proportion of responses to Intact probes (a), Related probes (b), Unrelated-Opposite probes (c), and Unrelated-Within probes (d) for Experiment 2. Lines at the top of the solid bars denote the group mean. Box around the mean indicates the 95% confidence interval of the mean. Jittered points denote individual participants' data. Note that splitting the Unrelated probes into two categories resulted in fewer trials per category, so the distribution of scores appears more clustered than for Intact and Related probes.

MPT results. We expanded the MPT model shown in Figure 3 to include two separate trees for Unrelated probes. For the Unrelated-Opposite probes, the tree is identical to the Unrelated tree in Figure 3. For the Unrelated-Within probes, an additional parameter (F) corresponds to retrieving a fuzzier level of representation, such that with probability F , a participant thinks the probe is similar to an original association (e.g., “this old man was somewhere in nature”) and then guesses whether the probe is “intact” (with probability a) or “related” (with probability $1-a$). With probability $1-F$, the participant does not retrieve a fuzzy level of representation, and the tree resolves to the Unrelated tree shown in Figure 3. We retained the equality restraint on the gist parameters (G_i and G_r) to make the model with the new F parameter overidentified. Tests of model fit with and without the equality restraint are reported in the supplement.

We collapsed the delay conditions within each age group because the 95% CIs of the difference in each parameter between the short and long delay blocks overlapped with 0. Table 1 provides a summary of parameter estimates, and Figure 4b depicts the difference score obtained by subtracting the posterior samples of the old from the young adults. These results are in line with Experiment 1. There were age differences on the verbatim memory parameters (V_i and V_r) but not the gist parameters (G_i and G_r). Older adults had a higher bias to respond “intact/related” to Unrelated probes (parameter b), which is likely attributable to the fact that they made more errors to Unrelated probes than the young adults did. Notably, the model fit the data well with the addition of a fuzzy retrieval parameter (F) in both age groups, which supports the notion that some participants may have retrieved a less specific level of representation for some of the Unrelated-Within probes. This parameter did not differ between groups.

We used the approach outlined in Experiment 1 to derive composite retrieval scores for the Intact and Related probes. Plugging in values from Table 1, composite retrieval scores for the Intact probes were .85 (95% CI: .78, .90) for the young adults and .70 (95% CI: .59, .78) for the old adults. For the Related probes, composite retrieval scores were .73 (95% CI: .62, .81) for the young adults and .61 (95% CI: .52, .69) for the old adults. There are two important aspects of these scores. First, older adults generally had lower overall retrieval scores than the young adults. Second, composite retrieval scores in Experiment 2 were generally lower than in Experiment 1. This suggests that when schematic support at encoding is reduced, participants retrieve less information about originally encoded associations, and this is especially true of older adults. Importantly, older adults exhibit the greatest deficit in retrieval of highly specific information (verbatim memory), show a less pronounced deficit in verbatim plus gist retrieval (composite memory score), and show relatively equivalent retrieval of fuzzy levels of specificity as the young adults (parameter F in the MPT model). In other words, the age-related deficits in associative memory observed in the present study scale with the amount of specificity needed-to-be-retrieved, consistent with the specificity principle of memory.

Discussion

Results from Experiment 2 were mostly consistent with and extended upon Experiment 1. Older adults exhibited the largest deficits in retrieval of associations at highly specific levels (e.g., “the old man was in *this* park”). They also demonstrated slightly lower composite retrieval (verbatim plus gist memory). Proportionally, they retrieved the gist of associations (e.g., “the old man was in *some* park”), upon failing to retrieve the specific representation, to similar extents as younger adults, but the pronounced loss of verbatim memory resulted in lower overall composite retrieval scores. However, they retrieved fuzzier levels of representation (e.g., “the old man was

outside somewhere”) to the same extent as younger adults. These findings point to a specificity gradient in associative episodic memory, whereby more specific information is harder to retrieve, especially with age, while representations are relatively more preserved at lower levels of specificity. This is consistent with Craik’s hierarchical representation model (Craik, 2002), but it extends on that model by offering firmer evidence for multiple levels of representation than has been shown in any previous study.

An important distinction between both experiments was the schematic support manipulation. When pairings were categorized (Experiment 1), composite retrieval scores were higher in both age groups than when this manipulation was absent (Experiment 2). The removal of categorized pairings made it possible that even the gist of an association might become susceptible to interference or forgetting. This appeared to occur to some extent in Experiment 2, for both young and older adults, a finding compatible with the specificity principle of memory (Surprenant & Neath, 2009).

As with Experiment 1, we found no delay effect on verbatim or gist memory. This supports the notion that difficulties with retrieving specific information occur on multiple time-scales, which is central to the universality of principles of memory (Surprenant & Neath, 2009). However, future research is warranted to assess whether specific or gist memory for associations declines with longer delays than those used in the present experiments.

General Discussion

Guided by Surprenant and Neath’s (2009) specificity principle of memory, we have shown that older adults are deficient, relative to younger adults, at retrieving associative episodic memories at highly specific levels of representation but can retrieve associations at lower levels of specificity (e.g., gist) to relatively similar extents as younger adults. To date, most of the

evidence suggesting that older adults rely on a gist-based processing strategy has come from studies employing categorized lists of material in which a semantic gist may be activated because all the items are related (see Devitt & Schacter, 2016; Tun, Wingfield, Rosen, & Blanchard, 1998). Our findings extend on these claims by showing that older adults can remember the gist of associations, which comprise the core of episodic memory (Tulving, 1983), providing perhaps the most direct evidence to date that older adults retrieve fuzzier episodic representations. Moreover, our findings of multiple levels of specificity for associations suggest that episodic memory representations might exist on a continuum (cf., Craik, 2002).

Although we suggest that older adults struggle to *retrieve* specific information, it is also possible they may fail to *encode* specific information (e.g., Naveh-Benjamin, 2000). However, Luo and Craik (2009) found that age-related deficits in memory for specific information are more likely attributable to retrieval than to encoding failures, but it is imperative for future research to continue to disentangle the effects of encoding versus retrieval specificity. In our view, it is likely a dynamic relationship between encoding and retrieval that influences older adults' memory for specific information, a view shared by Surprenant and Neath (2009). Further, as with any aging study employing a cross-sectional design, we cannot state that aging *causes* a loss of specific memory for associations as other factors may be at fault.

These results are consistent from both an MPT and signal detection framework. This is important given a growing debate as to whether recognition proceeds on a discrete or continuous basis (e.g., Pazzaglia et al., 2013). One important area in which this debate has not been extended much is the simplified conjoint recognition procedure (Stahl & Klauer, 2008), the type of paradigm used here. More work is also warranted in extending this debate to associative

recognition (e.g., Rotello, 2017), especially in conjunction with the measurement of verbatim and gist memory for associations.

We used face-scene pairs because they simulate remembering *where* people were encountered (e.g., Grupposo et al., 2007). Memory for where events occurred is one of the hallmarks of episodic memory (Tulving, 1983). The extent to which these findings would generalize to other types of associations in memory cannot be directly answered from the present study but would seem reasonable considering the similar decline they show with age (cf. Old & Naveh-Benjamin, 2008). For some associations, however, like remembering who committed a particular action, retrieval of the gist may not be sufficient, given potential legal implications for eyewitness testimony. Nevertheless, our findings that older adults can retrieve the gist of where they previously encountered people indicates that, at least for some associations in memory, age-related deficits are limited to retrieval of the specific representations.

Supplement to “A Specificity Principle of Memory: Evidence from Aging and Associative
Memory”

I. Hierarchical Bayesian Categorical Logistic Regression Model Diagnostics

Sampling quality. Model parameters were estimated from three independent MCMC chains, each with 2000 iterations (including 1000 warm up samples), for 6000 samples total, with convergence monitored by the $\hat{R}$ statistic. All $\hat{R} = 1.00$.

Posterior predictive checks: Density overlay plots comparing the fit of the observed data with 100 random datasets from the fitted model are depicted in Figures S1 and S2. All models provided adequate posterior predictive fits.

Fig. S1. Comparison of the observed data (y) to simulated data from the posterior predictive distribution (y_{rep}). Demonstrates satisfactory fit.

Fig. S2. Comparison of the observed data (y) to simulated data from the posterior predictive distribution (y_{rep}). Demonstrates satisfactory fit.

II. Multinomial Processing Tree (MPT) Model Diagnostics

Sampling routines. Model parameters were estimated from three independent MCMC chains, with an adaptation period of 15,000 samples, a burn in period of 10,000 samples, and a sampling period of 50,000 samples. All $\hat{R} \leq 1.03$.

Model fit. Model fit was assessed with the T_1 and T_2 statistics proposed in Klauer (2010), which measure discrepancies between the observed and predicted mean and covariance of individual frequencies, respectively. The number of samples was computed for which $T_{obs} < T_{pred}$, yielding a posterior predictive p value (PPP). Smaller values of PPP correspond to worse fit, and PPP close to 0 indicate insufficient fit.

The original MPT from the Simplified Conjoint Recognition paradigm (Stahl & Klauer, 2008) fit the data in both experiments poorly, all $PPP < .002$. The modified model reported in the main text fit the young ($T_{1obs} = 0.20$, $T_{1pred} = 0.19$, $PPP = 0.484$; $T_{2obs} = 47.16$, $T_{2pred} = 40.67$, $PPP = 0.388$) and old adults ($T_{1obs} = 0.17$, $T_{1pred} = 0.19$, $PPP = 0.556$; $T_{2obs} = 39.36$, $T_{2pred} = 41.92$, $PPP = 0.548$) well in Experiment 1. Additionally, the overidentified model with the additional fuzzy retrieval parameter reported in Experiment 2 fit both the young ($T_{1obs} = 0.25$, $T_{1pred} = 0.19$, $PPP = 0.300$; $T_{2obs} = 19.82$, $T_{2pred} = 16.31$, $PPP = 0.330$) and old adults ($T_{1obs} = 0.16$, $T_{1pred} = 0.18$, $PPP = 0.595$; $T_{2obs} = 22.21$, $T_{2pred} = 20.27$, $PPP = 0.423$) well, indicating that the addition of a fuzzy retrieval parameter (F) maintained satisfactory fit to the data. We also fit a less restricted version of the model in Experiment 2, by freeing the gist parameters. This model had satisfactory fit to the data, all $PPP \geq .294$, and yielded parameter estimates in line with those obtained with the overidentified model. Finally, setting F to 0 in the model resulted in substantial misfit for the T_1 statistic in the young adults, $T_{1obs} = 0.46$, $T_{1pred} = 0.19$, $PPP = 0.019$, and reduced fit for both

the T_1 and T_2 statistics among the old adults, both $PPP \leq 0.152$. Additional plots of model fit are available at <https://osf.io/9yige/>.

III. Signal Detection Theory (SDT) Analyses

Background. Traditionally, SDT models have been applied to tasks with a 2 x 2 matrix of responses, as in yes-no (both binary and ratings tasks) or forced choice recognition tasks. The simplified conjoint recognition procedure used in the present study expands the 2 x 2 response matrix to a 3 x 3 matrix (i.e., 3 probe types, each with 3 possible responses). Therefore, for each probe, there are two types of false alarms, each of which may have different psychological interpretations. A benefit of the MPT approach is that it readily accounts for different responses to each probe with parameters measuring psychological processes that are theoretically motivated (see Batchelder & Riefer, 1999). However, as mentioned in the main text, the discrete-state representation implied by the MPT approach has been the subject of intense debate in the recognition literature, so we took it upon ourselves to compute SDT estimates of discrimination and response bias for our design.

To get our 3 x 3 response matrix to the standard 2 x 2 format, we conducted a series of contrasts among the different probes. In one contrast, we compared Intact to Related probes, to derive a measure of high specificity discrimination. This contrast corresponds to the MPT parameters of verbatim memory because disentangling Intact from Related probes requires retrieving a highly specific representation of an originally encoded association. For this contrast, we retained only those trials where participants responded “intact” or “related” to Intact and Related probes (i.e., we dropped all Unrelated trials and all “unrelated” responses to Intact and Related probes). This was to ensure a singular type of hit (responding “intact” to Intact probes) and false alarm (responding “intact” to Related probes).

In a second contrast, we compared Intact plus Related probes to Unrelated probes. This contrast corresponds to the MPT approach of deriving a composite retrieval score because it combines information about how often participants retrieve associations at highly specific levels (e.g., correctly responding “intact” to an Intact probe, or “related” to a Related probe) with how often they retrieve associations at less specific levels of representation (e.g., erroneously responding “related” to an Intact probe, or “intact” to a Related probe).¹ We coded all “intact” and “related” responses to both Intact and Related probes as hits and all “unrelated” responses to these probes as false alarms.

Analyses. Although evidence from ROC curves typically favors the unequal variance model (see Rotello, 2017), we used the equal variance SDT model because the unequal variance model requires an additional parameter to account for the different variances of the signal and noise distributions. Typically, this parameter can be included by introducing more degrees of freedom with a ratings task. We included no ratings task, as confidence ratings can influence and even change the phenomena being studied (Brainerd, Nakamura, Reyna, & Holliday, 2017; Kellen, Winiger, Dunn, & Singman, 2018). Although we were limited to the equal variance model, it should be noted that the evidence in favor of the unequal variance model, or any SDT model as opposed to MPT models, is more mixed for associative recognition tasks, where receiver-operating characteristic curves are often more linear than those predicted by the unequal variance model (Rotello, 2017). Nevertheless, we acknowledge that the equal variance model has its limitations.

¹ One advantage of the MPT approach is that we can disentangle estimates of verbatim and gist memory from each other. However, there is no principled way to do this directly with the SDT model. Consider, for example, attempting a contrast focusing on only Intact versus Unrelated probes, and dropping all trials where participants respond “related.” Such a contrast excludes some gist information because responding “related” to Intact probes could be made on the basis of gist memory. This limitation motivated our approach to deriving composite retrieval scores from both the SDT and MPT models, but note that these scores reflect both verbatim and gist memory retrieval.

To derive estimates of discrimination, d' , and response bias, c for both contrasts, we used a hierarchical Bayesian SDT model (Rouder & Lu, 2005). The outcome on trial i for subject j is Bernoulli distributed, with probability p_{ij} that $y_{ij} = 1$.

$$y_{ij} \sim \text{Bernoulli}(p_{ij})$$

To map the probabilities to the real line, we used a probit link function, given by

$$\Phi(p_{ij}) = \beta_{0j} + \beta_{1j} * \text{Probe}_{ij}$$

where Φ is the cumulative normal density function. The intercept in the model corresponds to the standardized false alarm rate. To estimate c , we invert the sign of the intercept, i.e. $c = -\beta_{0j}$. The slope of the model corresponds to the increase in the hit rate for Intact probes (coded as 1) relative to recombined probes (coded as 0), and is thus our estimate of d' .

The subject-specific intercepts and slopes are drawn from a multivariate normal population-level distribution, with means μ_0 and μ_1 , and covariance matrix Σ ,

$$\begin{bmatrix} \beta_{0j} \\ \beta_{1j} \end{bmatrix} \sim \text{Normal} \left(\begin{bmatrix} \mu_0 \\ \mu_1 \end{bmatrix}, \Sigma \right)$$

Additional predictors can be added to the slope term. We included age (coded as 0 = young, 1 = old) as a between-subject predictor and delay as a within-subject predictor (coded as 0 = short, 1 = long). In Experiment 2, we also model the effect of type of unrelated probe (0 = opposite category, 1 = within category).

Models were estimated in the brms package for R (Bürkner, 2017). We specified weakly informative normal priors with means of 0 and standard deviations of 1 on the mean parameters, but retained the program's default priors on the variance and covariance parameters. The use of weakly informative priors is sufficient to regularize the posterior distribution and provides a principled method of correcting for perfect hit and false alarm rates, which otherwise result in incalculable estimates of d' and c (e.g., Stanislaw & Todorov, 1999). Model parameters were

estimated from three independent MCMC chains, each with 2000 iterations, with convergence monitored by the $\hat{R}$ statistic. All $\hat{R} = 1.00$.

Results - Experiment 1. Table S1 provides the marginal population-level estimates of d' and response bias c , split by age and delay. Higher values of d' correspond to greater sensitivity to detect associations as intact (for the Intact versus Related contrast) or intact/related (for the Intact plus Related versus Unrelated contrast). Positive values of c denote a more conservative bias (less likely to respond “intact”), whereas negative values of c correspond to a more liberal bias (more likely to respond “intact”).

For the Intact versus Related contrast, there was a main effect of age on both d' and c , with the older adults having lower discrimination, $\Delta d' = -0.85$ (95% CI: -1.16, -0.52), and a more liberal response bias than the young adults, $\Delta c = -0.61$ (95% CI: -0.82, -0.39), indicating they were more biased to respond “intact.” Regarding delay effects, the young adults had slightly lower estimates of d' in the long delay blocks, $\Delta d' = -0.28$ (95% CI: -0.52, -0.04), and slightly more liberal estimates of c , $\Delta c = -0.21$ (95% CI: -0.36, -0.07). For the old adults, there was no discernible effect of delay on discrimination, $\Delta d' = -0.14$ (95% CI: -0.36, 0.07), but a modest shift in response criterion, $\Delta c = -0.19$ (95% CI: -0.34, -0.04).

For the Intact + Related versus Unrelated contrast, there was no sufficient evidence for an age difference in discrimination, $\Delta d' = -0.30$ (95% CI: -0.83, 0.26), nor in response bias, $\Delta c = -0.13$ (95% CI: -0.53, 0.25). In addition, there was no discernible effect of delay on d' or c among the young or old adults, all 95% credible intervals overlapped with 0.

Table S1. SDT Estimates

Contrast	Young Adults				Old Adults			
	Short Delay		Long Delay		Short Delay		Long Delay	
	d'	c	d'	c	d'	c	d'	c
Intact vs Related (Exp 1)	1.39 [1.05, 1.73]	0.49 [0.29, 0.70]	1.08 [0.79, 1.38]	0.26 [0.13, 0.40]	0.41 [0.17, 0.65]	-0.18 [-0.36, 0.01]	0.28 [0.11, 0.44]	-0.35 [-0.53, -0.17]
Intact+Related vs Unrelated (Exp 1)	2.81 [2.34, 3.40]	1.13 [0.78, 1.50]	3.00 [2.50, 3.53]	1.30 [0.97, 1.66]	2.37 [1.96, 2.78]	0.96 [0.62, 1.29]	2.55 [2.19, 2.89]	1.02 [0.77, 1.25]
Intact vs Related (Exp 2)	1.50 [1.23, 1.77]	0.38 [0.21, 0.57]	1.77 [1.41, 2.17]	0.53 [0.36, 0.70]	0.75 [0.49, 1.02]	0.03 [-0.13, 0.19]	0.53 [0.31, 0.76]	-0.08 [-0.25, 0.09]
Intact+Related vs Unrelated-Opposite (Exp 2)	2.32 [1.88, 2.81]	1.33 [0.96, 1.76]	2.17 [1.78, 2.62]	1.20 [0.87, 1.59]	1.43 [1.21, 1.66]	0.66 [0.46, 0.87]	1.32 [1.11, 1.52]	0.55 [0.37, 0.74]
Intact+Related vs Unrelated Within (Exp 2)	1.76 [1.49, 2.06]	0.76 [0.57, 0.99]	1.65 [1.36, 1.96]	0.67 [0.46, 0.91]	1.20 [0.96, 1.44]	0.42 [0.21, 0.64]	1.08 [0.79, 1.38]	0.32 [0.06, 0.59]

Note. Values define the population-level mean and 95% Bayesian credible interval (in brackets).

MPT and SDT Comparisons – Experiment 1. The presence of an age effect on discrimination in the Intact versus Related contrast is consistent with the MPT estimated differences in verbatim retrieval. The absence of an age effect on discrimination in the Intact plus Related versus Unrelated contrast is consistent with the similar MPT estimates of composite retrieval for the young and old adults. The few delay effects reported above are consistent with the results from the logistic regression analyses, although the parameters of the MPT models were no different in the short and long delay blocks, and the delay effects on SDT estimates are relatively small. One notable distinction in the two modeling approaches is that the SDT results show that older adults had a more liberal bias in the Intact versus Related contrast, while the MPTs showed no differences in response bias parameters. This is likely attributable to the fact that in the MPT model, there are several processes that contribute to responding “intact,” including retrieving verbatim memory, guessing “intact” when only the gist is retrieved, and guessing “intact” when neither the verbatim nor the gist is retrieved.

Results - Experiment 2. Comparing the estimates of d' and c for Experiment 2 in Table S2 to those from Experiment 1, it is apparent that the Intact versus Related estimates are relatively similar in both experiments, but the Intact plus Related versus Unrelated estimates are somewhat lower than in Experiment 1. There are two logical explanations for this. First, we split the Unrelated probes into Unrelated-Opposite probes and Unrelated-Within probes, and the latter maintain fuzzier representations of the original associations. Second, by removing the schematic support manipulation (see the manuscript), we made it potentially harder for participants to remember even the gist of associations, such that the signal distribution shifts toward the noise distribution.

For the Intact versus Related contrast (assessing verbatim information), the results closely mimic those of Experiment 1. Older adults had poorer discrimination than the young adults, $\Delta d' = -0.85$ (95% CI: -1.16, -0.52), and had a more liberal response bias, $\Delta c = -0.46$ (95% CI: -0.65, -0.26). There was no discernible effect of delay on d' or c among the young or older adults, all 95% Bayesian credible intervals overlapped with 0.

The Intact plus Related versus Unrelated contrasts assessing composite retrieval demonstrated main effects of age for both the Unrelated-Within probes [$\Delta d' = -0.54$ (95% CI: -0.88, -0.19), and $\Delta c = -0.34$ (95% CI: -0.62, -0.06)] and the Unrelated-Opposite probes [$\Delta d' = -0.75$ (95% CI: -1.11, -0.35), and $\Delta c = -0.53$ (95% CI: -0.86, -0.22)]. This is consistent with the main effect of age reported in the logistic regression analyses. Importantly, within each age group, there was a main effect of type of Unrelated probe. For the young adults, $\Delta d' = -0.41$ (95% CI: -0.61, -0.18), and $\Delta c = -0.41$ (95% CI: -0.60, -0.20). For the old adults, $\Delta d' = -0.26$ (95% CI: -0.46, -0.07), and $\Delta c = -0.26$ (95% CI: -0.43, -0.09). In both age groups, participants had lower discrimination for the Unrelated-Within probes. There were no noteworthy delay effects on d' or c for either type of Unrelated probe among the young or older adults, all 95% CI's overlapped with 0.

MPT and SDT Comparisons – Experiment 2. The results from both modeling approaches are generally consistent. The age effect on discrimination in the Intact versus Related contrast in the SDT model corresponds to the age-related differences in verbatim parameter estimates in the MPT model. Further, the lower estimates of discrimination in the older adults for the Intact plus Related versus Unrelated contrast is consistent with the lower composite retrieval scores for the older adults obtained from the MPT approach. Also, the MPT estimates demonstrate that older adults had higher estimates of parameter b (corresponding to a bias to

respond “intact” or “related” even when the gist is not retrieved), which is consistent with the age-related differences in c in the Intact plus Related versus Unrelated contrasts. Importantly, in both age groups, discrimination scales with the amount of specificity that has to be retrieved. The lowest estimates of d' occur at the level of highest specificity, involved in discriminating Intact from Related probes. Estimates of d' are larger for discriminating Intact and Related probes from Unrelated-Within probes, which maintain some fuzzier degree of representation of original associations. Finally, estimates of d' are the largest for discriminating Intact and Related probes from Unrelated-Opposite probes. This provides additional support for levels of specificity.

References

- Batchelder, W. H., & Alexander, G. E. (2013). Discrete-state models: Comment on Pazzaglia, Dube, and Rotello (2013). *Psychological Bulletin*, 139, 1204-1212.
- Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, 6, 57-86.
- Boywitt, C. D., Kuhlmann, B. G., & Meiser, T. (2012). The role of source memory in older adults' recollective experience. *Psychology and Aging*, 27, 484-497.
- Brainerd, C. J., Gomes, C. F. A., & Moran, R. (2014). The two recollections. *Psychological Review*, 121, 563-599.
- Brainerd, C., Nakamura, K., Reyna, V., & Holliday, R. (2017). Overdistribution illusions: Categorical judgments produce them, confidence ratings reduce them. *Journal of Experimental Psychology: General*, 146, 20-40.
- Brainerd, C. J., & Reyna, V. F. (2015). Fuzzy-trace theory and lifespan cognitive development. *Developmental Review*, 38, 89-121.
- Bürkner, P. C. (2017). Brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80, 1-28.
- Castel, A. D. (2005). Memory for grocery prices in younger and older adults: The role of schematic support. *Psychology and Aging*, 20, 718-721.
- Castel, A. D., McGillivray, S., & Worden, K. M. (2013). Back to the future: Past and future era-based schematic support and associative memory for prices in younger and older adults. *Psychology and Aging*, 28, 996-1003.
- Chen, T., & Naveh-Benjamin, M. (2012). Assessing the associative deficit of older adults in long-term and short-term/working memory. *Psychology and Aging*, 27, 666-682.

- Craik, F. I. M. (2002). Levels of processing: Past, present...and future? *Memory*, 10, 305-318.
- Devitt, A. L., & Schacter, D. L. (2016). False memories with age: Neural and cognitive underpinnings. *Neuropsychologia*, 91, 346-359.
- Dixon, P. (2008). Models of accuracy in repeated-measures designs. *Journal of Memory and Language*, 59, 447-456.
- Dube, C., & Rotello, C. M. (2012). Binary ROCs in perception and recognition memory are curved. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 38, 130-151.
- Ebner, N. C., Riediger, M., & Lindenberger, U. (2010). FACES –a database of facial expressions in young, middle-aged, and older women and men: Development and validation. *Behavioral Research Methods*, 42, 351-362.
- Gelman, A., Jakulin, A., Pittau, M. G., & Su, Y. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2, 1360-1383.
- Gruppuso, V., Lindsay, D. S., & Masson, M. E. J. (2007). I'd know that face anywhere! *Psychonomic Bulletin & Review*, 14, 1085-1089.
- Heck, D. W., Arnold, N. R., & Arnold, D. (2018). TreeBUGS: An R package for hierarchical multinomial-processing-tree modeling. *Behavioral Research Methods*, 50, 264-284.
- Kellen, D., Winiger, S., Dunn, J. C., & Singmann, H. (2018, August 30). Testing the Foundations of Signal Detection Theory in Recognition Memory.
<https://doi.org/10.31234/osf.io/p5rj9>
- Klauer, K. C. (2010). Hierarchical multinomial processing tree models: A latent-trait approach. *Psychometrika*, 75, 70-98.

- Konkle, T., Brady, T. F., Alvarez, G. A. and Oliva, A. (2010). Scene memory is more detailed than you think: the role of categories in visual long-term memory. *Psychological Science*, 21, 1551-1556.
- Kruschke, J. K. (2011). Bayesian assessment of null values via parameter estimation and model comparison. *Perspectives on Psychological Science*, 6, 299-312.
- Kruschke, J. K. (2018). Rejecting or accepting parameter values in Bayesian estimation. *Advances in Methods and Practices in Psychological Science*, 1, 270-280.
- Luo, L., & Craik, F. I. M. (2009). Age differences in recollection: Specificity effects at retrieval. *Journal of Memory and Language*, 60, 421-436.
- Naveh-Benjamin, M. (2000). Adult age differences in memory performance: Tests of an associative deficit hypothesis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 26, 1170-1187.
- Old, S. R., & Naveh-Benjamin, M. (2008). Differential effects of age on item and associative measures of memory: A meta-analysis. *Psychology and Aging*, 23, 104-118.
- Pazzaglia, A. M., Dube, C., & Rotello, C. M. (2013). A critical comparison of discrete-state and continuous models of recognition memory: Implications for recognition and beyond. *Psychological Bulletin*, 139, 1173-1203.
- R Core Team. (2018). R: A language and environment for statistical computing. R version 3.4.4 (2018-03-15) [Computer software]. Vienna, Austria: Author. Retrieved from <http://www.R-project.org/>
- Reyna, V. F., & Brainerd, C. J. (1991). Fuzzy-trace theory and framing effects in choice: Gist extraction, truncation, and conversion. *Journal of Behavioral Decision Making*, 4, 249-262.

- Reyna, V. F., & Brainerd, C. J. (1995). Fuzzy-trace theory: An interim synthesis. *Learning and Individual Differences*, 7, 1-75.
- Roediger, H. L. III (2008). Relativity of remembering: Why the laws of memory vanished. *Annual Review of Psychology*, 59, 225-254.
- Roediger, H. L. III, & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 21, 803-814.
- Rotello, C. M. (2017). Signal detection theories of recognition memory. In J. T. Wixted (Ed.), *Learning and Memory: A Comprehensive Reference*, 2nd edition (Vol. 4: *Cognitive Psychology of Memory*). Elsevier.
- Rouder, J. N. (2014). Optional stopping: No problem for Bayesians. *Psychonomic Bulletin & Review*, 21, 301-308.
- Rouder, J. N., & Lu, J. (2005). An introduction to Bayesian hierarchical models with an application in the theory of signal detection. *Psychonomic Bulletin & Review*, 12, 573-604.
- Smith, J. B., & Batchelder, W. H. (2008). Assessing individual differences in categorical data. *Psychonomic Bulletin & Review*, 15, 713-731.
- Smith, J. B., & Batchelder, W. H. (2010). Beta-MPT: Multinomial processing tree models for addressing individual differences. *Journal of Mathematical Psychology*, 54, 167-183.
- Stahl, C., & Klauer, K. C. (2008). A simplified conjoint recognition paradigm for the measurement of gist and verbatim memory. *Journal of Experimental Psychology: Learning, Memory, & cognition*, 34, 570-586.

- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, 31, 137-149.
- Surprenant, A. M., & Neath, I. (2009). *Principles of Memory*. New York, NY: Taylor & Francis Group, Psychology Press
- Tulving, E. (1983). *Elements of Episodic Memory*. Oxford, UK: Clarendon Press.
- Tun, P. A., Wingfield, A., Rosen, M. J., & Blanchard, L. (1998). Response latencies for false memories: Gist-based processes in normal aging. *Psychology and Aging*, 13, 230-241.