Paradigmatic formation through context-mediation

Kevin D. Shabahang (k.shabahang@gmail.com)
School of Psychological Sciences, The University of Melbourne, Australia
Hyungwook Yim (hwyim@hanyang.ac.kr)
Department of Cognitive Sciences, Hanyang University, Seoul, Republic of Korea
School of Psychological Sciences, The University of Melbourne, Australia
Simon J. Dennis (simon.dennis@unimelb.edu.au)
School of Psychological Sciences, The University of Melbourne, Australia

Abstract

Words that regularly fill the same sentential slots are said to be paradigmatically related. Paradigmatic relations may be retained through a direct association or a latent representation at encoding, or by reinstating context during retrieval. We paired proper names by embedding them into two instances of the same sentence frame, each in a separate list, yielding blocks of two study-cloze sessions. The pairing between proper names was fixed across twelve blocks. In the static condition, the same sentence frames were used across blocks, while in the dynamic condition sentence frames changed for each block. Interference should accrue in both conditions if paradigmatic relations are based on a direct association or overlap in a latent representation, however, if paradigmatic relations are mediated by retrieved context then changing the sentence frame should release interference. Our results are consistent with a context-mediation account of paradigmatic relations.

Keywords: Context reinstatement; Paradigmatic relations; Interference; Language and thought; Learning

Take the two sentences “there is coffee in the mug”, “there is tea in the mug. Words that occur in the same context like COFFEE and MUG are said to have a first-order, or syntagmatic association. Words that appear in the same slots (but different sentences) are said to have a second order, or paradigmatic association. Syntagmatic associations determine what words are to be bound together in the same context. Paradigmatic associations align words based on the extent to which they appear in the same position relative to other words and cluster according to semantic roles.

Despite its theoretical importance as an organizing principle, empirical support for the psychological reality of paradigmatic relations has been inconclusive. Anderson and Bower (1971) carried out some experiments by extending early work by Rohwer and Lynch (1967). Rohwer and Lynch varied the number of times a verb was repeated in a list of 16 Subject-Verb-Object sentences and had participants try to recall the objects when cued with the subjects, after studying all the sentences. Surprisingly there was no effect of verb repetition on recall rates. Anderson and Bower (1971) replicated Rohwer and Lynch’s null finding when analyzing accuracy, but an analysis of intrusions by same-verb distractors showed a monotonically increasing function of verb repetition. They attributed the result to a guessing strategy, but the pattern could indicate the formation of paradigmatic relations between subjects and objects.

Interpolated learning paradigms incorporating sentence learning can be generalized as A-B, A-C designs, where participants are first given a set of A-B type paired associates (e.g., “The dog chased Eric”) for study, followed by an additional study list. Some of the pairs in the second list will contain the left-hand item of a pair in the first list, but this time paired with another item (A-C or “The dog chased Joe”). After studying the A-C pairs, followed by a potential test, participants are cued with the A items (“The dog chased”) and asked to recall the corresponding item based on the pairings in the first list (B or Eric). Memory for the original pairs in the A-B, A-C condition is compared against a control (A-B, C-D) to determine the level of interference or facilitation. Both pre-experimental and experimental factors determine the probability of facilitation relative to interference.

One pre-experimental factor is the similarity between right-hand items (B and C), or *arguments*. The occurrence of A-C following the presentation of A-B is more likely to facilitate the later retrieval of the original A-B pair, the more similar the changed argument (i.e., C in A-C) is to the original argument (i.e., B in A-B). On the other hand, decreasing similarity between the two arguments increases the likelihood of interference when retrieving A-B pairs (Osgood, 1949). Mensinck and Raaijmakers (1988) explain the relationship between argument similarity and interference in terms of pre-experimental contextual overlap. When the A-C pair is presented, the high degree of shared contextual overlap between C and B leads to the retrieval of A-B and the bound contextual information, resulting in the simultaneous activation of A-B and A-C and re-encoding as a unitary item-to-context bundle.

Anderson and Bower (1973) had participants go through five study-test phases of a list of sentences, followed by six interpolated study-test phases of three other sentence lists

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

(two study-test phases per list). Finally, participants were cued with a sentence frame containing the object, and were to respond with the corresponding subject noun and verb. Anderson and Bower varied the degree of overlap between sentences in the interpolated lists and sentences in the original list with respect to their logical and grammatical alignment. The sentence "Nixon kicked the baby" is maximally aligned with "Nixon kicked the milkman", because Nixon is both the agent and the grammatical subject in both sentences. An active-to-passive (or vice versa) transformation inverts grammatical alignment as in "The milkman was kicked by Nixon", but preserves logical alignment since Nixon's role remains the agent. Logical alignment is removed by downgrading Nixon to the patient role, again either with intact grammatical alignment "Nixon was kicked by the milkman" or without "The milkman kicked Nixon". Anderson and Bower found that cued recall for the original list degraded when a constituent in the original sentence (e.g., Nixon) appeared in the interpolated lists, regardless of the quality of alignment according to the logical/grammatical factorization. Critically, despite high overall interference, there was no evidence for paradigmatic interference (i.e., same-role). That is, the occurrence of Nixon as the agent in an interpolated list led to the same level of interference as the occurrence of Nixon as the patient.

A similar pattern of results was presented in a few later studies. Anderson (1975) used both cued-recall and recognition and were unable to find paradigmatic interference. In their cued-recall experiment, they found decreasing accuracy when a target concept was repeated in an interpolated list sentence and in their recognition experiment, they found increasing response-times in verification (or rejection) of target concepts. The pattern in both experiments occurred irrespective of the alignment between logical and grammatical role of target and interpolated sentences.

Dosher (1983) found some evidence for broad role-level representations by having participants read Subject-Verb-Object (SVO) sentences, followed by a later true-false recognition test. Dosher included lures with different levels of alignment with the studied sentences in grammatical role as in earlier studies by Anderson and Bower, but the semantic role representations were not verb-specific and simply differentiated agents from patients. Participants were slower to correctly reject lures made by swapping either the subject or noun of one studied sentence with the subject or noun of another, when swapped constituents both played the same broad semantic role in the studied sentences. The pattern was reliable for two out of three types of lures. For the third type of lure, the pattern was present only if both study and test sentences were grammatically aligned – both active or both passive voice – but reversed otherwise. In another experiment, Dosher cued participants with verbs corresponding to some studied sentences and found that participants were more likely to confuse the agents of sentences for one-another and the patients with one-another than confusing the agents with the patients. Intrusions matched the correct words in terms of broad semantic role.

Demonstrating the formation of paradigmatic relations through experimentally manipulated stimuli has been evasive, particularly when accuracy is used as the main dependent measure. There have been glimpses of their formation upon the analysis of error however. For example, Anderson and Bower's (1971) finding of increasing same-verb object intrusions as a function of verb repetition when participants were cued with the subject alone is indicative of paradigmatic formation. Repeating the verbs across study sentences likely facilitated the alignment of words that appeared as the subjects of the same verbs. Likewise, words that appeared as objects would paradigmatically align. Cueing with a given subject would be more confusable with other paradigmatically aligned subjects, therefore increasing the probability of producing the incorrect object. On the other hand, the correct object's paradigmatic alignment with other incorrect objects would increase the probability of a same-verb intrusion.

One reason stronger evidence for paradigmatic relations has been difficult to obtain is that interpolated learning tasks simultaneously elicit facilitation and interference. For example, Aue, Criss, and Novak (2017) showed proactive facilitation instead of inhibition in a set of cued-recall tasks. They had participants learn a set of face-word paired-associates and later had them recall the paired word when presented with the corresponding face. In a later list, participants were given faces from the first list, but paired with new words. When later cued with a face, participants were better at responding to faces that also occurred in the first list relative to those that did not. Their results are most consistent with the view that the second presentation of the face sometimes cues the word with which it was paired in the first list.

Experimental factors like the retention interval between the presentation of the interpolated pairs (i.e., A-C) and the final memory test for the A-B pairs and the lag between the A-B and A-C pairs both influence whether memory of A-B is facilitated or inhibited. Bruce and Weaver (1973) show that including A-C pairs in between presentation and test of A-B pairs can sometimes facilitate the later retrieval of A-B pairs to the same extent as if the A-B pairs were presented twice. Robbins and Bray (1974a; 1974b) varied the lag between the original A-B pairs and the overlapping A-C pairs in addition to the retention interval between the interpolated pairs and final test for the A-B pairs. They found facilitation when the lag between the original pairs and the overlapping pairs was small (around 5 to 10 seconds) and the retention interval was long (25-50 seconds).

Jacoby, Wahlheim and Kelley (2015) asked participants if a paired associate changed (i.e., A-B in list one and A-C in list two) during the presentation of an interpolated list and found facilitation for A-B pairs when change in A-C pairs was recollected and inhibition otherwise. Notably the A-C intrusions remained reliably higher than control even when change was noticed and despite facilitation, suggesting that performance is driven through the simultaneous interaction between facilitation and interference. Recent data by Frankland and Greene (2020) provides consistent neurophysiological support for the duality. They presented participants with a list of sentences to read as their brains were scanned and found that participants that were given sentences of the form “the hawk surprised the moose” and “the hawk surprised the cow” – which have an A-B, A-C structure – representational analysis of their functional brain imaging data indicated that the pattern of neural activity driving the encoding of the B and C items is correlated in the anterior-medial prefrontal cortex (amPFC) but anti-correlated in a hippocampal sub-region of interest.

Jacoby et al. (2015) explained facilitation of A-B pairs through A-C pairs by assuming that the A-C pairs sometimes cue the retrieval of the A-B pairs, therefore leading to the further strengthening of the A-B pairs. They also show that the probability of facilitation is much higher when the A-B and A-C pairs occur within a short temporal distance of one another. As in Mensinck and Raaijmackers (1988), Jacoby et al. explain the dependence on lag in terms of contextual overlap. Since the contexts surrounding the study of A-B and A-C are likely to have more overlap the closer they are in time, then lower lags increase the probability that A-C’s shared context with A-B will trigger the latter’s retrieval.

The tendency for A-C pairs to cue the retrieval of A-B pairs may be a mechanism for the formation of paradigmatic associations. For example, if the sentence “The dog bit Eric” and “The dog bit Joe” were used as the A-B, and A-C items, respectively, then presenting “The dog bit Joe” may evoke the retrieval of “The dog bit Eric”, and hence the formation of an associative link between Eric and Joe due to their simultaneous activation in working memory. Alternatively, a history of shared context without a direct association may be sufficient for the formation of paradigmatic relations. Yet a third possibility is that paradigmatic relations form in a latent/hidden representation through the continuous prediction of contexts from items or items from contexts (e.g., Mikolov et al. 2013).

Here, we focus on differentiating the three possible mechanisms with which paradigmatic relations may form. Based on previous results demonstrating difficulties finding a reliable behavioural signature for the formation of experimentally manipulated paradigmatic relations, we adopt a task that attempts to maximize interference between subsets of studied sentences and yields a large number of datapoints for the exploration of patterns of intrusions during recall. Given that an A-B pair was studied, and after some delay, an A-C pair was studied, C and B can be called paradigmatically overlapping due to their shared context, A. The three competing accounts make different assumptions about how a paradigmatic relation can form between C and B.

The **direct-association account** assumes that the presentation of the A-C pair may cue the retrieval of the earlier A-B pair, leading to the strengthening of an association between B and C. On the other hand, suppose A-B and A-C are encoded, then cueing with either B (or C) should activate A, and then A can activate C (or B). Alternatively, cueing with A should activate both B and C. The **context-mediation account** assumes that that context drives the formation of paradigmatic relations independently of direct associations through its reinstatement. The **latent-representation account** assumes that shared context is a catalyst for the formation of similar hidden unit patterns through representational alignment in a latent space.

The direct-association and latent-representation accounts both assume that paradigmatic relations form during encoding whereas the context-mediation account assumes that they form during retrieval. The direct-association and latent-representation accounts imply that once a paradigmatic relation has formed, it is relatively independent of the corresponding overlapping context(s), but they differ in that the direct-association account predicts a faster rate of paradigmatic formation relative to the latent-representation account which relies on slower error-driven learning processes. The mediated-context account can be distinguished from the other two accounts because disrupting the contextual reinstatement process should prevent paradigmatic interference.

Suppose that each block contained two study-test phases. In the first study-test phase, participants read A-B pairs, and are later cued with A, and asked to produce the B items. In the next study-test phase, participants read A-C pairs, and are cued with A items for retrieval of the C items. The direct-association and context-mediation accounts predict that B items will interfere with C items in the second phase of the first block. The two accounts make different predictions if the following block contains D-B and D-C pairs in the two phases. The direct-association account predicts further build-up of interference in both phases since the direct association between B and C strengthens every time they share a context, but the context-mediation account assumes release from interference in the first phase and build-up of interference in the second phase. The latent-representation account does not predict interference until the A-B and A-C pairs have been repeated over several trials.

Method

Participants

We recruited 149 (52 female and 97 male) through Amazon’s Mechanical Turk. The majority of participants were between 31 to 40 years of age, and reported to be native English speakers.

Materials

We grouped 294 different transitive verbs into 156 sets, ensuring that verbs in different sets were as unrelated as possible. Some sets contained a single verb (e.g., acknowledge) whereas others contained multiple (e.g., arrest, apprehend, and bust), ranging from 1 to 11 verbs. For each verb (e.g., cure), we used the lemma names corresponding to the hyponyms of a relevant Wordnet (Miller, 1995) synset (e.g., medicine.n.02) to obtain a base set of common nouns (e.g., tonic and restorative) to be used as subject arguments. We discarded low frequency verbs and nouns and odd verb-noun pairings in each of the 156 sets. The number of nouns in each verb set ranged from 2 to 139 (M = 60.54; SD = 51.96) in the remaining pool. Finally, 45 common English proper names, each ranging between 4 to 7 characters (M = 4.73; SD = 0.93), were also selected to be used as targets.

The common nouns were used as subjects and the proper nouns were used as objects in SVO sentences. For each participant, depending on the condition, either 12 (Static), 72 (Dynamic), or 144 (Random) sentence frames were sampled from the pool with the constraint that only one verb can be used from each verb set and there are no repetitions in the common nouns. We sampled an additional 12 sentences, with no overlap in verb-set, common noun, or proper name for each condition to be used as practice trials. We constructed a cloze item for each of the sentences by replacing the proper noun (i.e., the object) with an underscore.

Procedure

For each participant, twelve proper names were sampled and used as targets in all 24 study-cloze phases. We grouped the study-cloze phases into 12 blocks, each with two consecutive study-cloze phases. Six of the 12 names were randomly assigned to the first (PhaseA) and second lists (PhaseB) in each block.

Table 1 shows two example study-cloze items for the Static, Dynamic, and Random groups, across phases (i.e., lists) in the first two blocks. The first two lists show are from the first block and the third and fourth lists are from the second block. The same two sentence frames (italicized) repeat across phase 1 and phase 2 lists for both the Static and Dynamic groups, but only in the first block. In the second block, new frames are introduced in the Dynamic group but the same ones are repeated again in the Static group. New frames are introduced according to the same structure in each additional block for the Dynamic group while the same frames are repeated throughout blocks for the static group. Stimuli corresponding to the Random group are constructed by ensuring that all sentence frames are unique across lists.

We used the JsPsych library (de Leeuw, 2015) for conducting the experiment. Each participant initiated the experiment through Amazon's Mechanical Turk, and was randomly assigned to one of the three context groups upon giving informed consent. Instructions were presented stating that "in this study, you will be shown separate lists of six sentences. After you see each list, you will be given fill-in-the-blank questions about that list. You should **answer based on the list you just studied** and not any list before. You will be presented with twenty four separate lists." Three follow-up multiple-choice questions ensured that the participant understood the instructions and helped defend against automated online agents. Participants were allowed three attempts to correctly answer all three questions. Otherwise the experiment was terminated. Upon completing the comprehension check, participants were asked demographic questions and proceeded to the practice trials.

Table 1: Example study and cloze frames.

List	Static	Dynamic	Random
1	*The dog chased* Eric	*The dog chased* Eric	*The dog chased* Eric
	The taxi drove Mary	*The taxi drove* Mary	*The taxi drove* Mary
	The dog chased _	*The dog chased* _	*The dog chased* _
	The taxi drove _	*The taxi drove* _	*The taxi drove* _
2	*The dog chased* Joe	*The dog chased* Joe	*The truck hit* Joe
	The taxi drove Lucy	*The taxi drove* Lucy	*The looker dated* Lucy
	The dog chased _	*The dog chased* _	*The truck hit* _
	The taxi drove _	*The taxi drove* _	*The looker dated* _
3	*The dog chased* Eric	*The remedy healed* Eric	*The remedy healed* Eric
	The taxi drove Mary	*The skater excited* Mary	*The skater excited* Mary
	The dog chased _	*The remedy healed* _	*The remedy healed* _
	The taxi drove _	*The skater excited* _	*The skater excited* _
4	*The dog chased* Joe	*The remedy healed* Joe	*The boss hired* Joe
	The taxi drove Lucy	*The skater excited* Lucy	*The cop watched* Lucy
	The dog chased _	*The remedy healed* _	*The boss hired* _
	The taxi drove _	*The skater excited* _	*The cop watched* _

Note. The lists in bold-face correspond to the second block.

Practice started by reminding participants to "answer based on the list you just studied and not any lists before", followed by a single block of two study-cloze phases. The structure of the practice sentence frames corresponded to the participants' group allocation. For each study-cloze phase, six sentences were presented in the center of the screen, one-at-a-time, replaced with a fixation cross between each sentence. Each sentence was presented for 5 seconds, with an inter-sentence stimulus interval of 1 second. A filler task followed after the participant was presented with the sixth sentence in the list. In the filler task, two random numbers between 1 and 100 were presented to the participant, and they were prompted to answer if the sum was odd or even. After 5 filler trials, participants were presented with each of the corresponding cloze items, one-at-a-time, in addition to a text-box where they were to provide the target (i.e., the proper name). The participants were allowed to leave the text-box blank and proceed to the next cloze item. They were provided feedback about their accuracy after each cloze item during the practice trials, however, no feedback was provided in the main experiment. Following practice, participants initiated the main 12 blocks with a keypress after being informed that they are entering the main session. We randomized the order of the sentences and cloze items for each of the 24 study and cloze lists.

The cycle of context shifts is the key variable in the present experiment. Contexts cycle in each study-cloze phase in the Random group, imposing no systematic overlap between targets across phases. At the other extreme, there is a complete absence of context shifts in the Static group. In between the two extremes, context shifts in each block for the Dynamic condition, imposing overlap between contexts across each phase of each block. Critically, the link between targets across phases remains the same in all blocks in the Dynamic group. If paradigmatic relations correspond to the formation of direct links between associates (e.g., Eric-Henry or Mary-Lucy), then the probability with which a subject confuses the two should increase across blocks in both the Static and Dynamic context groups. On the other hand, if paradigmatic relations form through context reinstatement during retrieval, then shifting context should facilitate release from interference in the Dynamic group, but continue to rise in the Static group. For the Random group, the probability that a target in the first phase is confused with a target in the second phase should be at random. In the first phase of the first block, (Table 1) the targets are Eric and Mary, but in the second phase, the targets are Joe and Lucy. For each cloze test, we define the filler as the word that was paired with the same sentence frame in the prior phase. With the exception of the very first phase, when no prior names are paired in the experiment, targets in the first phase are always fillers in the second phase, and vice-versa. We operationalize paradigmatic confusion as the probability with which a filler is reported, given that one of the twelve study names were reported. That is, we define the probability of a filler by discarding all extralist trials (c.f. Anderson & Bower, 1972).

The direct association account predicts that, when names are repeatedly paired with the same contexts, then the probability with which a filler is reported will increase rapidly in both the Static and Dynamic groups, relative to the Random group. The context mediation account predicts that the probability of reporting a filler will increase across blocks in the Static group, but rise and fall across phases, within a block, in the Dynamic group. The latent representation account predicts that interference will increase gradually across blocks, regardless of phase.

Results

We computed the probability each participant correctly answered a cloze test correctly separately for each list as a measure of accuracy. Additionally, given an intralist intrusion (i.e., one of the other names in the list) we computed the probability with which the intrusion was a filler (i.e., paradigmatic associate) as a measure of paradigmatic interference.

Figure 1 shows accuracy as a function of list, separately for each group. A 3 x 24 Bayesian Mixed ANOVA with one group factor (Static, Dynamic, Random) and one within-subjects factor (List 1 to List 24) did not indicate any reliable difference in accuracy across lists and groups. The main effects of group and list and the interaction term had Bayes factors below 0.15.

Figure 2 shows the probability of reporting a filler name instead of the target, between groups and across lists. As would be expected, the probability that two random names across lists in a block are confused is near zero. At the other extreme, the probability of a filler increases across the first four lists in the Static group, and remains high through the middle lists, with a slight decrease near the end lists. The probability of filler intrusions shifts between low to high in the Dynamic condition, showing a zig-zag pattern across lists.

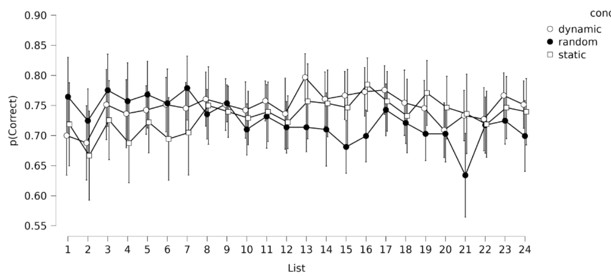

Figure 1: Probability of correct recall.

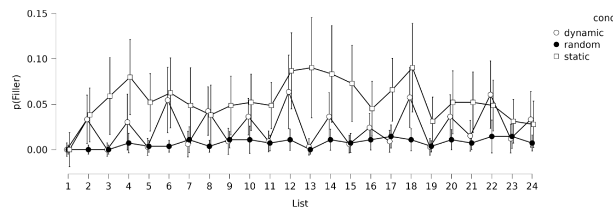

Figure 2: Probability of paradigmatic intrusion.

We averaged filler probabilities across each of the 12 odd- and even-numbered lists to obtain the probability of filler for each block phase separately. A 3 x 2 Bayesian Mixed ANOVA with one group factor (Static, Dynamic, Random) and one within-subjects factor (Odd, Even) favoured the model with an interaction between phase and group in addition to both main effects (Bayes factor = 283913244). Four pairwise comparisons further showed that the probability of filler was higher in the Static group relative to the Random group, for both the Odd (Bayes factor = 124.234) and Even (Bayes factor = 5052.572) lists. In contrast, the probability of filler for the Dynamic group was higher than in the Random group, for Even lists (Bayes factor = 498.862), but not Odd lists (Bayes factor = 0.212).

Discussion

We proposed three competing mechanisms that may underlie the formation of paradigmatic overlap between words, setting two encoding accounts (direct-association and latent-representation) against a pure retrieval account (context-mediation). We manipulated contextual overlap between target items to distinguish between the accounts. The target items were paradigmatically aligned through repeated pairing either through the same contexts in the static group (A-B, A-C, A-B, A-C, …) or different contexts in the dynamic group (A-B, A-C, D-B, D-C, …) and compared interference relative to a control group (A-B, C-D, E-B, G-D, …). We argued that both the direct-association and latent-representation account assume growing interference as more lists are studied in the static and dynamic group through a decoupling of source context and paradigmatic relationship formation. Alternatively, we argued that if the locus of paradigmatic formation is at retrieval process and always mediated through shared context, then interference should change depending on whether it is the first list in a block (i.e., when the context just shifted) or the second list.

Our results support the context-mediation account. Interference, as indicated by the probability of a filler, increased steadily when contextual overlap was static, however, it showed a zig-zag pattern indicative of the reduction in interference each time the context shifted and a build-up of interference when it was repeated with another set of targets. When filler probabilities were averaged based on the study phase, the probability of filler in the first phase of the dynamic group was no different from the random group while filler probabilities in both phases of the static group were greater than the random group even after 12 repetitions.

The context-mediation account of paradigmatic formation is consistent with results on the different time-course of retrieving associative versus relational information. Dosher (1984) had participants study word pairs in a paired-associate recognition task that were either semantically related (e.g., PURSUE-FOLLOW) or unrelated (e.g., OPEN-VEGETABLE). At test, participants were presented with lures that were either semantically related or unrelated, and were trained to respond as fast as possible upon an auditory signal. Dosher systematically varied the latency between the response signal and the presentation of each probe to capture different snapshots of the temporal profile of each recognition judgment. Results showed that participants were more likely to falsely recognize a semantically related lure (e.g., STIFLE-SUPPRESS) relative to an unrelated lure (e.g., FIRE-STIFLE) early on during processing (within around one second) but the rates become identical shortly after (after about two seconds). They explained the results based on a late-suppression model, where cueing with a pair initially reinstates all contexts associated with the items, but is later restricted to exclude pre-experimental context. In a similar response signal design, Ratcliff and McKoon (1989) constructed stimuli in the same manner as Anderson and Bower (1973), factorially combining grammatical and logical alignment in a sentence learning true-false recognition memory experiment. Their results showed an initially high propensity towards endorsing lures that were grammatically aligned with the target ("Nixon kicked the milkman") but were in logical misalignment ("Nixon was kicked by the milkman"), followed by the suppression of the response later during processing.

The idea that paradigmatic relations form online through the retrieval of relevant contextual information can explain the pattern observed in Dosher (1984) and Ratcliff and McKoon (1989). When a cue like "Nixon was kicked by the milkman" is presented as a probe, the item-to-context and item-to-item associations globally activate associated contexts and items in a single bottom-up pass. The contextual activations then constrain earlier actications through a later top-down stream, connecting contexts-to-items.

In conclusion, our results are consistent with the view that the formation of paradigmation relations occurs online through the reinstatement of shared context.

References

Anderson, J. R., Bower, G. H. (1971). On an associative trace for sentence memory. *Journal of Verbal Learning and Verbal Behaviour, 10*, 673-680.

Anderson, J. R., & Bower, G. H. (1973). Human Associative Memory.

Anderson, J. R. (1975). Item-specific and relation-specific interference in sentence memory. *Journal of Experimental Psychology: Human Learning and Memory, 1(3)*, 249.

Aue, W. R., Criss, A. H., Novak, M. D. (2017). Evaluating mechanisms of proactive facilitation in cued recall. *Journal of Memory and Language, 94*, 103-118.

Bruce, D., & Weaver, G. E. (1973). Retroactive facilitation in short-term retention of minimally learned paired associates. *Journal of Experimental Psychology, 100(1)*, 9.

De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods, 47(1)*, 1-12.

Dosher, B. A. (1983). Relational information in sentence memory. *Journal of Psycholinguistic Research, 12(4)*, 377-396.

Dosher, B. A. (1984). Discriminating preexperimental (semantic) from learned (episodic) associations: A speed-accuracy study. *Cognitive psychology, 16(4)*, 519-555.

Frankland, S. M., & Greene, J. D. (2020). Two ways to build a thought: distinct forms of compositional semantic representation across brain regions. *Cerebral Cortex, 30(6)*, 3838-3855.

George A. Miller (1995). WordNet: A Lexical Database for English. *Communications of the ACM Vol. 38, No. 11*: 39-41.

Jacoby, L. L., Wahlheim, C. N., & Kelley, C. M. (2015). Memory consequences of looking back to notice change: Retroactive and proactive facilitation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 41(5)*, 1282.

Mensink, G. J., & Raaijmakers, J. G. (1988). A model for interference and forgetting. *Psychological Review, 95(4)*, 434.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.

Osgood, C. E. (1949). The similarity paradox in human learning: A resolution. *Psychological review, 56(3)*, 132.

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology, 21(2)*, 139-155.

Robbins, D., & Bray, J. F. (1974a). Repetition effects and retroactive facilitation: Immediate and delayed recall performance. *Bulletin of the Psychonomic Society, 3(5)*, 347-349.

Robbins, D., & Bray, J. F. (1974). The spacing effect and the AB, AC paradigm: Evidence for retroactive facilitation. *Journal of Experimental Psychology, 103(3)*, 420.

Rohwer Jr, W. D., & Lynch, S. (1967). Form class and intralist similarity in paired-associate learning. *Journal of Verbal Learning and Verbal Behavior, 6(4)*, 551-554.