



Orthographic priming from unrelated primes: Heterogeneous feedforward inhibition predicted by associative learning

James S. Adelman^{a,*}, Iliyana V. Trifonova^{a,b,c}

^a Department of Psychology, University of Warwick, UK

^b Department of Psychology, University of York, UK

^c School of Psychology, University of East Anglia, UK

ARTICLE INFO

Keywords:

Orthographic processing
Lexical decision
Masked priming
Associative learning

ABSTRACT

A common assumption among models of orthographic processing is that letter-word inhibitory relationships all share the same strength: activity in the letter B has the same impact on a word like RACE as does equivalent activity in the letter F. However, basic associative learning mechanisms imply that the existence of the neighbor word FACE gives more opportunity to learn a negative weight from the neighbor letter F than from the non-neighbor letter B, leading to stronger negative letter-word weights for neighbor than non-neighbor letters. In masked primed lexical decision, therefore, *fity*, a *neighborly* prime formed using neighbor letters, should be a more inhibitory prime for RACE than *bund* (vice versa for LARK). We present simulations of weight learning using Rescorla and Wagner's (1972) equations and three experiments consistent with this prediction. Further simulations show heterogeneous feedforward connections from letters to words could contribute to phenomena previously attributed to lexical competition.

Introduction

There's no I in TEAM, but that there's no R in TEAM is more important for the reader who must discriminate the words TEAR, TERM, TRAM and REAM from TEAM. Models of orthographic processing do not, however, distinguish between the relationship between I and TEAM and that between R and TEAM. Instead, the assumption is made that mismatching letter information is equally discriminatory, regardless of the wrong letter identity.

A canonical example of this assumption comes in the interactive-activation model (McClelland & Rumelhart, 1981): In this model, a single letter-word inhibitory connection weight parameter describes the extent to which the activity of letter nodes leads to inhibitory input to mismatching words. Although interactions from other model components complicate predictions, the *direct* effect of a wrong letter is the same for all 25 (in English) possible wrong letters in this model, and in every model (e.g., Adelman, 2011; Davis, 2010; Grainger & van Heuven, 2003; Norris, 2006) accounting for effects in orthographic priming paradigms commonly used to study lexical selection, a similar equivalency holds even though the mechanisms differ.

The most common orthographic priming paradigm is form-primed lexical decision (Forster, Davis, Schoknecht, & Carter, 1987). On each

trial, a mask is presented before a brief (ca. 50 ms) *prime* lower-case letter string, which is followed by an upper-case *target* letter string; the participant indicates whether the target is a real word or not. On trials with word targets and nonword primes, latencies are shorter when many letters are shared between prime and target than when only few are (see Adelman et al., 2014, for a review). As such, a typical control condition in these experiments has primes with no letters in common with the corresponding targets. Although there is evidence that pronounceability or orthographic legality of control primes may affect latencies, the identity of letters in control primes as they relate to the targets is not normally considered further, as they are universally theoretically expected to have homogeneous effect.

From the perspective that the discriminations among words must be learned, however, such homogeneity is far from expected. A learner has more opportunity to learn that F is negative evidence for RACE – from experiencing FACE, an orthographic *neighbor* (differing in a single letter; Coltheart, Davelaar, Jonasson, & Besner, 1977) – than that learner does to learn that B does not occur in RACE. Discriminating compound cues with error-driven learning leads to negative association weights (Rescorla & Wagner, 1972) because the positive prediction of the shared part of the compound (ACE) is erroneous on non-consequential (FACE) trials and must be counteracted by the uniquely non-consequential part

* Corresponding author at: Department of Psychology, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, UK.

E-mail address: J.S.Adelman@warwick.ac.uk (J.S. Adelman).

<https://doi.org/10.1016/j.jml.2022.104372>

Received 8 July 2022; Received in revised form 22 August 2022; Accepted 22 August 2022

Available online 9 September 2022

0749-596X/© 2022 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

(F is present when RACE is absent). Without the positive prediction to counteract, when B occurs in items dissimilar to RACE, no learning of a negative association should occur.

Said another way: Model weights learned in this way (i.e., by Rescorla & Wagner's, 1972, rules and related formulations) reflect differences in conditional probability of the recipient word (conditional contrasts; Shanks, 1995) between cases with and without the source letter, in each of the relevant contexts. For a letter not in the word, the only such non-zero contrast is that for the context of other letters (_ACE): When the neighbor letter is present (F in FACE), the probability of the recipient word (RACE) is zero, but when the neighbor letter is absent, the probability of the recipient word is positive, but not necessarily one (because there also exist LACE and MACE). The contrast between these is negative, so when averaged together with the irrelevant zero contrasts, a negative weight – conditioned inhibition – results.

The consequence for priming is that a *neighborly* prime constructed from neighbor letters, such as fity-RACE, should have an inhibitory effect compared to one constructed from non-neighbor letters, such as bund-RACE. We present modeling that demonstrates that the relevant negative weights are learned by the Rescorla and Wagner (1972) learning rules, and two experiments that test the prediction for these types of stimuli. We then present further simulations and another experiment based on six-letter targets, and primes in which only two letters are manipulated.

Data availability

Simulation code (C++) and results, DMDX experimental files, experimental data, and R code for experimental data analysis are available from OSF <https://osf.io/htcqn/> or from <https://adelmanlab.org/neighborly>.

Simulations 1A-1D: Neighborly primes, four-letter stimuli

For the purposes of illustration, we trained four simple error-driven learning models with vocabularies of four-letter words with different sublexical orthographic representations. The purpose of training four models was to demonstrate that the prediction is not a product of a specific encoding scheme. The models make the same qualitative predictions.

Method

Letter-level representations

A: Slot-coded model Each letter had a distinct representational unit for each of the four ordinal positions within a four-letter word. When a stimulus was presented to the model, these were *on* if the letter was present in that position in the stimulus, and *off* otherwise. An additional unit that was always *on* was expected to learn a frequency-related bias. This is the simplest method of coding position of letters as in the McClelland and Rumelhart (1981) model.

B: Position-free letters model Each letter had a set of representational units that encoded the number, but not the position of the letters. One of these units was *on* whenever the letter occurred at least once in the stimulus; another was *on* whenever the letter occurred at least twice in the stimulus; and so on. In all other cases, the units were *off*. An additional (bias) unit was always *on*. This was included to show that the predicted result does not rely on any positional encoding.

C: Position-free letters and bigrams model As the position-free letters model, except additionally, each possible open bigram (defined as ordered pairs of [not necessarily adjacent] letters in a string – for instance, WORD contains open bigrams W_O, W_R, W_D, O_R, O_D, R_D) had six units similarly encoding number. This was included as the preceding encoding was not able to distinguish anagrams; this dual encoding of letters and bigrams is used in more recent iterations of open bigram models (Snell, van Leipsig, Grainger, & Meeter, 2018).

D: Sloppy slots As the slot-coded model, except that when a letter was presented in a position p , the nodes for that letter identity in all positions q were activated to a level $e^{-(q-p)^2}$, so that the correct position was *on* with activity 1, and activity for nearby letter nodes for the same identity decreases as distance increases.

Vocabulary and word-level representations

All four-letter words in SUBTLEX-UK (van Heuven, Mandera, Keuleers, & Brysbaert, 2014) had a representational unit in the model that was connected to every letter-level unit with a to-be-learned weight. Its *response* to any stimulus was the sum of weights connecting it to units that were turned *on* by the stimulus. (This would be equivalent to the net input in an interactive activation model.) That is, for a word i , the response $r_i = \sum l_j w_{ij}$ where $l_j = 1$ for letters (or bigrams) j present in the stimulus, and $l_j = 0$ for other letter-level units, except for the sloppy slots model, where l_j was the activity level described above.

Training

All weights were initiated at zero. Twenty million training trials were run, drawing words to be presented from the vocabulary randomly with probability in proportion to their SUBTLEX-UK frequency. The relevant letter-level representations were turned *on*, and the response in all the word units was calculated. The target t_i for a word i was one for the presented word, and zero for all other words. A proportion of the discrepancy (error) between target and response is added to the weights that produced the response. That is, $w_{ij} \leftarrow w_{ij} + \alpha l_j (t_i - r_i)$ for all words i and for all letters j , except that overlearning was not treated as error; no change was made if $r_i > t_i = 1$ or $r_i < t_i = 0$. (This primarily ensures that negative weights learned where discrimination is important – such as neighborly letter links learned on trials with neighbors – are not pushed back to zero on irrelevant learning trials where the prediction is negative rather than zero.) The simulations presented here use a learning rate of $\alpha = 0.002$. Note that weights involving absent letters are not updated (because $l_j = 0$ for those letters), but both present and absent words have corresponding weights updated.

Testing stimuli

We selected 348 four-letter words that had a neighbor in each position (from among 639 in our initial list constructed from a UNIX word list¹; this was 27% of the four-letter words) so that a (unique) nonword *neighborly* prime could be constructed, formed of letters that when substituted for another letter in the target another word was formed. No attempt was made to control pronounceability, but most primes were pronounceable, as the manner of construction tended to produce primes whose consonant-vowel structure was the same as the target. Words were paired such that the neighborly prime for each word of the pair contained no letters that could form a word from the paired word by a single substitution and so was a neutral control prime for the paired word (e.g., bund is the neighborly prime of LARK, and the control prime of RACE). Thus, the neighborly and control primes were not different stimuli, they were just paired with different targets to create the different conditions. Trial and error suggested that 174 was the maximal number of pairs of unique words that we were able to construct within this constraint and the requirement of the later experiments that the number be a multiple of 3 for counterbalancing purposes.

Results

We presented each of the prime stimuli and measured the response in the corresponding target word units. Table 1 shows that the models produced more negative target responses from neighborly primes than

¹ These stimuli were chosen from the UNIX word list before the simulations were trained using SUBTLEX-UK. All the chosen word targets appeared in SUBTLEX-UK.

Table 1

Target responses to each prime type for the four models for the four-letter stimuli in Simulations 1A-1D.

Model	Prime type		
	Identity	Control	Neighborly
1A: Slot-coded	0.490	-0.663	-1.127
1B: Letters	0.366	-0.798	-0.892
1C: Letters and bigrams	0.702	-0.301	-0.410
1D: Sloppy Slots	0.477	-0.622	-1.077

from controls, whereas identity primes produced facilitatory responses. We interpret this prediction only ordinally, because the function that links these activations to response times should be nonlinear (according to any plausible model of how responses are made) and so other numerical properties (differences, ratios) of these predictions about activation would not be borne out in response times. We also do not interpret differences between models because (a) of nonlinearity between activations and response times and (b) the Rescorla-Wagner formulation allows more parameter variation than we have explored and so any differences between model predictions are not fixed.

Experiment 1

The first experiment tested the prediction in the standard primed lexical decision paradigm.

Method

Participants

After exclusion of non-native speakers (17), further participants with accuracy below 75% (2), participants excess to counterbalancing (2), and those whose data was lost to equipment failure (2), data of 54 undergraduate students receiving partial course credit² were available for analysis.

Stimuli

We used the word targets and corresponding primes that were test stimuli for the simulations. We constructed a nonword target from each word target (in most cases, by altering only one letter) for which a (unique) neighborly prime could be constructed, and further constructed a control nonword prime for these nonword targets; these were not balanced in the same way as the word primes as we had no hypothesis of interest regarding nonword latencies.

Design

Lexical decision latencies to word targets were measured following neighborly, neutral control and identity primes from all participants. Six counterbalancing lists of 348 trials were constructed. First, two target lists were constructed containing only one word of each pair, and the nonword constructed from its partner (i.e., participants either saw RACE and LIRK, or RAPE and LARK). From each, three lists were constructed, each with an equal number of word and nonword targets associated with each prime type, so that each prime-target combination occurred in exactly one list.

Procedure

Participants were instructed that their task was to examine the upper-case letter string on each trial and press the right shift key if it was

² This sample size was determined as follows: All participants available to us through this course credit scheme in a single academic year were allocated to one of the two experiments presented here to produce equal final Ns. This, as anticipated, yielded samples substantially larger than had been needed to observe the effect in pilot work.

a real English word or the left shift key if it was not. Following 12 practice trials, all trials from the relevant counterbalancing list were presented in a new random order.

On each trial, DMDX (Forster & Forster, 2003) displayed the ##### mask for 500 ms, then the 12.5pt lower-case prime for 50 ms, then the 20pt upper-case target until response or a maximum of 3000 ms. Feedback was given after incorrect responses only.

Results

Correct word responses with latencies between 150 and 1500 ms whose targets received at least 60% correct responses were analyzed. Mean latencies are displayed in Table 2. The latencies were analyzed with a linear mixed effect model for the fixed effect of prime type; the model with full random slopes of prime type on participant had singular fit³, so only random intercepts for participant were included. The omnibus ANOVA showed the prime types to differ significantly, $\chi^2(2) = 83.09, p < .001$. All pairwise comparisons were significant; as predicted, neighborly primes yielded significantly longer latencies than control primes, $\chi^2(1) = 6.58, p = .010$.

Experiment 2

To reduce a potential effect of temporal merging of prime and target – that the f of fity and the ACE of RACE merge to form an illusory FACE – the second experiment introduced a mask between prime and target.

Method

Participants

After exclusion of non-native speakers of English (19), data for 54 undergraduate participants receiving partial course credit were available.

Stimuli, design and procedure

As Experiment 1, except a 30 ms presentation of %%%% occurred between prime and target.

Results

Analysis proceeded as for Experiment 1; mean correct word latencies are displayed in Table 2. Models with random slopes produced singular fits so only random intercepts were included⁴. The omnibus ANOVA showed the prime types to differ significantly, $\chi^2(2) = 13.63, p = .001$.

Table 2

Mean correct word target response times (ms) in Experiments 1–3.

Experiment	Prime type		
	Identity	Control	Neighborly
<i>Four-letter stimuli</i>			
1: #####-prime-TARGET	630	656	667
2: #####-prime-%%%%-TARGET	642	646	655
<i>Six-letter stimuli</i>			
3: #####-prime-TARGET	692	717	727

³ There was no significant evidence that the missing random slopes yielded superior fit, $\chi^2(5) = 2.67, p = .751$. Other models we investigated led to the same substantive conclusion regarding the comparison of neighborly and control conditions, irrespective of the singular fit.

⁴ The singular model with random slopes did not have a significantly superior fit, $\chi^2(10) = 8.29, p = .060$. Other models we investigated led to the same substantive conclusion regarding the comparison of neighborly and control conditions, irrespective of the singular fit.

Neighborly prime latencies significantly differed from the other prime types; as predicted, neighborly primes yielded significantly longer latencies than control primes, $\chi^2(1) = 4.70, p = .030$.

Discussion of experiments 1 & 2

Experiments 1 and 2 both provided evidence in line with the prediction that primes composed solely of non-target letters from neighbors of the target would lead to longer response times than primes composed solely of letters not appearing in any neighbor. All the targets in these experiments, due to the manner of stimulus construction, had a neighbor in all (four) positions. Although having neighbors across all positions is reasonably common among four-letter words, it is rare more generally. This, and the use of the same stimuli in both experiments, can lead to a concern regarding the generality of the findings. In the following, therefore, we examine six-letter targets, and manipulate only two letters of the prime.

Simulations 2A-D: Neighborly primes, six-letter stimuli

Method

All simulations were run in the same way as Simulations 1A-D, except that (a) the number of units was increased to accommodate six-letter words; (b) six-letter training stimuli were selected from SUBTLEX-UK (van Heuven et al., 2014) that were listed with frequency of at least 50 and were present in either or both of the UK and US spell-check; and (c) new six-letter testing stimuli were used, selected as described below.

Testing stimuli

We selected 180 six-letter words that had neighbors in at least two positions. For each of these words, we constructed a neighborly prime that included two of the letters that if substituted into the target would produce a neighbor word, in the positions that they would do so, with the four other positions filled with letters that were not in any neighbor. Words were paired so that the neighborly prime for each word was a control prime containing no neighbor letters for the other word, and the non-neighbor letters were the same in the paired primes. For instance, fxvzbx-PACING [facing, paving] and jxszbx-POKING [joking, posing] were neighborly primes and targets in one pair. The non-neighbor letters were chosen randomly from those not appearing in targets or their neighbors, resulting in very few pronounceable primes.

Results

As for the four-letter stimuli, the simulations with six-letter stimuli showed more negative targets responses to the neighborly primes than to the control primes, as can be seen in Table 3.

Experiment 3

The prediction for the six-letter stimuli was thus the same as that for the four-letter stimuli and was tested in the next experiment.

Table 3
Target responses to each prime type for the four models for the six-letter stimuli in Simulations 2A-2D.

Model	Prime type		
	Identity	Control	Neighborly
2A: Slot-coded	0.516	-0.419	-0.742
2B: Letters	0.404	-0.637	-0.766
2C: Letters and bigrams	0.852	-0.214	-0.285
2D: Sloppy Slots	0.567	-0.513	-0.874

Method

Participants

After exclusions of non-native speakers (36), participants with accuracy less than 75% (4) and the last-collected participants in some counterbalancing lists to equate numbers across those lists (3), data from 84 participants were available for analysis. We sought (and obtained) a larger number of participants⁵ than the preceding experiments due to concerns regarding power: The manipulation was weaker because it affected two out of six letters rather than four out of four.

Stimuli

The word stimuli from Simulations 2A-D were augmented with an equal number of nonwords to act as lexical decision foils.

Design

Lexical decision latencies were measured following neighborly, identity and control primes. Four counterbalancing lists were created, all of which contained every target once. Each member of a word target pairing was primed by the identity prime in two lists, and its partner in the other two lists. In one of the lists where the target was not primed with the identity prime, it was primed by the neighborly prime, and in the last, it was primed by the control prime (its partner's neighborly prime). This ensured that similar primes (the two neighborly primes from the same pair) did not appear in the same list and ensured half of trials were identity primed. We sought to increase the proportion of identity primes (from a third in the preceding experiments to a half) because of concerns regarding power: The manipulation was weaker, and we believe that a high incidence of related primes increases priming (cf. Bodner & Masson, 2001).

Procedure

As Experiment 1, except the mask was #####.

Results

Analysis proceeded as for the prior experiments. Mean correct word target response times are shown in Table 2. Models with random slopes (for either or both of participants and items) returned singular fits, so we report results from the model with random intercepts for participants and items⁶. The omnibus ANOVA showed the prime types to differ significantly, $\chi^2(2) = 105.77, p < .001$. All pairwise comparisons were significant; as predicted, neighborly primes yielded significantly longer latencies than control primes, $\chi^2(1) = 5.56, p = .018$.

Discussion

Experiment 3 replicated the finding from Experiments 1 and 2 that primes containing non-target neighbor letters produced longer latencies than primes that contained only non-target non-neighbor letters. It generalized the result from primes that had four out of four such neighbor letters to primes that had only two out of six such neighbor

⁵ This was achieved by using all participants available to us in a year, rather than splitting them over two related experiments, as was done for Experiments 1 and 2.

⁶ There was, however, evidence that random slopes improved the fit of the model, $\chi^2(10) = 10.64, p = .032$. Other models we investigated led to the same substantive conclusion regarding the comparison of neighborly and control conditions, irrespective of the singular fit. Models leading to the same substantive conclusion included one with only random slopes for the identity vs. control contrast. This model did not differ from the full random slopes model, $\chi^2(6) = 0.94, p = .988$. For this model, the omnibus comparison was $\chi^2(2) = 65.58, p < .001$ and for the comparison of neighbourly and control primes $\chi^2(1) = 5.61, p = .018$.

letters, to longer targets with fewer neighbors generally, and to primes that were unpronounceable rather than primes that were pronounceable.

General discussion

In all the experiments, primes constructed using non-target neighbor letters led to longer latencies than primes constructed entirely of other non-target letters. This effect was predicted from simple associative learning theory because the reader has relevant experience that these neighbor letters predict the absence of the target word when other cues suggest it might be present. Interposing a mask between the prime and target might have removed the effect if it were due to a simple form of visual blending, but the effect remained.

The effect does not follow directly from letter-word links in current models of orthographic processing. Moreover, there is little reason to believe it could emerge from other interactive processes in models: For a substantial effect of this type to arise as a result of lateral inhibition, at least one relevant neighbor word would need to be activated by the single constituent letter and remain active despite inhibitory influences from (a) the three or five letters of the prime not in that neighbor and (b) other words activated by the letters of the prime. An alternative explanation in which the neighbor is activated by some combination of prime and target letters is not compatible with the rapid changes in letter activation due to stimulus changes normally assumed in models, and in any case is not compatible with the result that the observed effects is preserved when the prime and target are separated by a mask⁷. The present experiments therefore provide strong evidence for heterogeneous connection strengths from the letter level to the word level. This finding could lead to some reinterpretation of several previously reported effects in the orthographic processing literature. These phenomena previously attributed to lexical inhibition may be linked, at least in part, to heterogeneous feedforward connections, as would be predicted from associative learning. Such associative learning mechanisms are implicated throughout cognition and have been previously investigated in a variety of psycholinguistic contexts (e.g., Baayen et al., 2011; Ramscar et al., 2010). However, even where orthographic learning has been included in larger models using associative rules (e.g., Hendrix, Ramscar & Baayen, 2019), the implications for orthographic processing have not been fully explored.

Associative learning and shared neighbor primes

Error-driven associative learning also predicts that positive associations will differ in their strength; these associations would be weaker for elements that are shared with a neighbor than those that are not. The links from L, A and Y to LAZY will be weakened on learning trials when LADY is learned, but not those from Z; in contrast, there are no learning trials on which Z → LAZY is strongly unlearned and A → LAZY is not. Therefore, a prime for LAZY that contains L, A and Y will be weaker than one containing L, Z and Y. van Heuven, Dijkstra, Grainger, and Schrievers (2001) derived the same prediction regarding prime strength from lateral inhibition in an interactive activation model and tested and confirmed it in an experiment with Dutch four-letter words comparing of shared neighbor primes like *laby*-LAZY, no-shared neighbor primes like

⁷ The Spatial Coding Model (Davis, 2010) which has such inhibitory mechanisms qualitatively produces the effect, but the magnitude of the effect is 0.2 cycles, which is considered equivalent to 0.2 ms, for Experiment 1, and less for Experiment 2. One might be tempted to infer that this simply means the model's inhibitory parameters should be increased but note that Trifonova and Adelman's (2018) sandwich priming results suggest that the inhibitory influences in this model are already too strong.

lozy-LAZY and control primes.

Simulations 3A-3D⁸ with Van Heuven et al.'s (2001) stimuli – whose results are in Table 4 – confirmed that the pattern arises as described from associative learning. Thus, van Heuven et al.'s finding that shared neighbor primes like *laby*-LAZY produce less facilitatory priming than no-shared neighbor primes like *lozy*-LAZY, might not occur because items like *laby* activate competitors like LADY which inhibit targets like LAZY, but because the neighbor-position Z → LAZY connection involving a unique letter is stronger than the non-neighbor-position A → LAZY connection involving a shared letter.

Associative learning and the prime lexicality effect

Other previous priming experiments that have been taken in support of the existence of lexical competition (lateral inhibition at the word level) naturally used stimuli that differ in the learnability of the wrong letters in the primes. Comparisons of word neighbor primes with nonword neighbor primes cannot be done without comparing letters with and without relevant learning opportunities. By definition, a word neighbor prime like *axle* for ABLE contains a neighborly letter (x) that should be strongly negatively associated with the target, whereas a nonword neighbor prime does not. We investigated the consequences for the effect of prime lexicality (in interaction with prime relatedness) in the models trained for Simulations 1A-1D, with the stimuli of Davis and Lupker (2006)'s Experiment 1, comparing priming effects of word and nonword neighbor primes, relative to control primes, on word targets.

The results in Table 5 show that while responses are similar for word and nonword control primes, word neighbor primes yielded a weaker response than nonword neighbor primes. Although Davis and Lupker (2006) found a reversal of the priming effect for word primes, and therefore these simulations only partially explain the observed pattern, Forster and Veres (1998) have found that the magnitude of this interaction can be affected by strategic factors (based on the type of nonword foil) so explaining this effect fully would always rely on the role of decision-making mechanisms and not just lower-level connection weights.

Methodological implication of the finding

Methodologically, this finding also emphasizes the importance of fully reporting all primes for orthographic priming effects in lexical decision involving supposedly unrelated letters; this includes the assignment of each specific control prime for each target when primes are re-used across conditions.

Conclusion

In sum, we have provided evidence for heterogeneous feedforward

Table 4

Target responses to each prime type for the four models for van Heuven et al.'s (2001) Dutch stimuli in Simulations 3A-3D.

Model	Prime type		
	No shared neighbor	Shared neighbor	Control
3A: Slot-coded	0.152	0.097	-0.247
3B: Letters	0.041	0.035	-0.199
3C: Letters and bigrams	0.192	0.129	-0.197
3D: Sloppy Slots	0.139	0.102	-0.333

⁸ Given the language of the stimuli, we trained a new set of models with four-letter words from SUBTLEX-NL (Keuleers, Brysbaert, & New, 2010) that occurred in more than 2 documents in the same manner as Simulations 1A-D.

Table 5

Target responses to each prime type for the four models for Davis & Lupker's (2006) Experiment 1 stimuli in Simulations 4A-4D.

Prime relatedness: Model	Prime lexicality			
	Nonword		Word	
	Neighbor	Control	Neighbor	Control
4A: Slot-coded	0.408	-0.725	0.123	-0.701
4B: Letters	0.114	-1.110	-0.014	-1.124
4C: Letters and bigrams	0.260	-0.351	-0.016	-0.343
4D: Sloppy Slots	0.336	-0.753	0.091	-0.732

inhibition through empirical differences in efficacy between two prime types both of which would previously have been considered unrelated primes of equivalent quality. We anticipated such differences as a generalization of conditioned inhibition in associative learning and simulations based on associative learning theory predicted the qualitative pattern of our experiments. Similar simulations of other inhibitory phenomena in orthographic priming suggest that inhibitory patterns should not be automatically attributed to online lateral inhibition. Our experimental results suggest that there is variation in the strength of inhibitory negative letter-word associations. Moreover, our simulations of these results produce variation in the strength of both positive and negative associations that predict other experimental patterns that have been attributed to lexical competition in the form of lateral inhibition. Nevertheless, these simulations do not constitute complete models of the priming process or visual word recognition, so it remains to be seen how effectively a broader model incorporating such heterogeneous connections can be compatible with a broader range of phenomena. At present, it also remains to be seen whether any other theoretical account can reasonably model the neighborly prime inhibition that has been empirically demonstrated here.

Authorship

JSA conceived the idea, ran the simulations, selected stimuli, analyzed the data, programmed and co-ordinated the running of Experiment 3, and drafted the manuscript and its revision. IVT programmed and ran Experiments 1 and 2 and checked the data analysis of all experiments. Both authors revised the draft manuscripts.

Note

We thank David Harvey for conducting pilot work; Inchara Athreya, Phoebe Exelby, Fizz Hunter, Fatima Hussain Kanwal, Davit Kim and Jessica Theo-Chinwo for assistance with Experiment 3; and Zachary Estes and Gordon Brown for helpful discussions.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

CRediT authorship contribution statement

James S. Adelman: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Supervision, Validation, Writing – original draft, Writing – review & editing. **Iliyana V. Trifonova:** Investigation, Software, Validation, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Adelman, J. S. (2011). Letters in Time and Retinotopic Space. *Psychological Review*, 118, 570–582.
- Adelman, J. S., Johnson, R. L., McCormick, S. F., McKague, M., Kinoshita, S., Bowers, J. S., ... Davis, C. J. (2014). A behavioral database for masked form priming. *Behavior Research Methods*, 46, 1052–1067.
- Baayen, R. H., Milin, P., Filipovic Durdevic, D., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, 118, 438–482.
- Bodner, G. E., & Masson, M. E. J. (2001). Prime validity affects masked repetition priming: Evidence for an episodic resource account of priming. *Journal of Memory & Language*, 45, 616–647.
- Coltheart, M., Davelaar, E., Jonasson, J., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and performance VI* (pp. 535–555). Erlbaum.
- Davis, C. J. (2010). The spatial coding model of visual word identification. *Psychological Review*, 117, 713–758. <https://doi.org/10.1037/a0019738>
- Davis, C. J., & Lupker, S. J. (2006). Masked inhibitory priming in English: Evidence for lexical inhibition. *Journal of Experimental Psychology: Human Perception and Performance*, 32, 668–687.
- Forster, K. I., & Forster, J. C. (2003). DMDX: A Windows display program with millisecond accuracy. *Behavior Research Methods, Instruments, & Computers*, 35(1), 116–124.
- Forster, K. I., Davis, C., Schoknecht, C., & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *The Quarterly Journal of Experimental Psychology*, 39A, 211–251. <https://doi.org/10.1080/14640748708401785>
- Forster, K. I., & Veres, C. (1998). The prime lexicality effect: Form-priming as a function of prime awareness, lexical status, and discrimination difficulty. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 498–514.
- Grainger, J., & van Heuven, W. J. B. (2003). Modeling letter position coding in printed word perception. In P. Bonin (Ed.), *The mental lexicon* (pp. 1–23). Nova Science.
- Hendrix, P., Ramscar, M., & Baayen, R. H. (2019). NDRA: A single route model of response times in the reading aloud task based on discriminative learning. *PLoS ONE*, 14(7), e0218802.
- Keuleers, E., Brysbaert, M., & New, B. (2010). SUBTLEX-NL: A new frequency measure for Dutch words based on film subtitles. *Behavior Research Methods*, 42, 643–650.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: An account of basic findings. *Psychological Review*, 88, 375–407. <https://doi.org/10.1037/0033-295X.88.5.375>
- Norris, D. (2006). The Bayesian reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological Review*, 113, 327–357. <https://doi.org/10.1037/0033-295X.113.2.327>
- Ramscar, M., Yarlett, D., Dye, M., Denny, K., & Thorpe, K. (2010). The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science*, 34(6), 909–957.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black, & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). Appleton-Century-Crofts.
- Shanks, D. R. (1995). Is Human Learning Rational? *The Quarterly Journal of Experimental Psychology*, 48A, 257–279. <https://doi.org/10.1080/14640749508401390>
- Snell, J., van Leipsig, S., Grainger, J., & Meeter, M. (2018). OB1-reader: A model of word recognition and eye movements in text reading. *Psychological Review*, 125, 969–984.
- Trifonova, I. V., & Adelman, J. S. (2018). The sandwich priming paradigm does not reduce lexical competitor effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44, 1743–1764.
- van Heuven, W. J. B., Dijkstra, T., Grainger, J., & Schriefers, H. (2001). Shared neighborhood effects in masked orthographic priming. *Psychonomic Bulletin & Review*, 8, 96–101.
- van Heuven, W. J. B., Mandera, P., Keuleers, E., & Brysbaert, M. (2014). SUBTLEX-UK: A new and improved word frequency database for British English. *The Quarterly Journal of Experimental Psychology*, 67, 1176–1190. <https://doi.org/10.1080/17470218.2013.850521>