

The Isolation, Primacy, and Recency Effects Predicted by an Adaptive LTD/LTP Threshold in Postsynaptic Cells

Sverker Sikström

Lund University Cognitive Science (LUCS), Lund University

Received 26 July 2004; received in revised form 10 August 2005; accepted 16 August 2005

Abstract

An item that stands out (is isolated) from its context is better remembered than an item consistent with the context. This isolation effect cannot be accounted for by increased attention, because it occurs when the isolated item is presented as the first item, or by impoverished memory of nonisolated items, because the isolated item is better remembered than a control list consisting of equally different items. The isolation effect is seldom experimentally or theoretically related to the primacy or the recency effects—that is, the improved performance on the first few and last items, respectively, on the serial position curve. The primacy effect cannot easily be accounted for by rehearsal in short-term memory because it occurs when rehearsal is eliminated. This article suggests that the primacy, the recency, and the isolation effects can be accounted for by experience-dependent synaptic plasticity in neural cells. Neurological empirical data suggest that the threshold that determines whether cells will show long-term potentiation (LTP) or long-term depression (LTD) varies as a function of recent postsynaptic activity and that synaptic plasticity is bounded. By implementing an adaptive LTP–LTD threshold in an artificial neural network, the various aspects of the isolation, the primacy, and the recency effects are accounted for, whereas none of these phenomena are accounted for if the threshold is constant. This theory suggests a possible link between the cognitive and the neurological levels.

Keywords: Isolation effect; Primacy; Recency; Adaptive LTP–LTD threshold; Synaptic plasticity; Neural network; Model

1. Introduction

It is well known that events that are different from the prevailing context are better recalled in free recall tests than events that are congruent with the context (Schmidt, 1991). That is, isolated items, or items that are incongruent with the surrounding context, are better recalled than nonisolated items consistent with the current context, a phenomenon known as the *isolation*

Correspondence should be addressed to Sverker Sikström, Lund University Cognitive Science (LUCS), Kungshuset, Lundagård, S-222 22 Lund, Sweden. E-mail: sverker.sikstrom@lucs.lu.se

effect. For example, a letter that is presented in a list that otherwise consists of numbers is better recalled than the numbers. Traditional accounts of the isolation effect have emphasized cognitive concepts such as distinctiveness and salience (Green, 1956; Jenkins & Postman, 1948). The isolation effect is a nontrivial phenomenon. As argued later, it cannot easily be accounted for by mechanisms such as salience, distinctiveness, or interference. The *primacy effect* is the finding that the free recall performance is improved for the first few items (Murdock, 1962). Primacy cannot easily, as reviewed later, be accounted for purely by rehearsal from short-term memory. The *recency effect* is the finding that the last few items are better recalled on immediate testing (Murdock, 1962). The isolation, primacy, and recency effects share the property that they occur when there is a change in a given context; however, these phenomena are seldom theoretically or experimentally related. A goal of this article is to establish a common theoretical framework.

The purpose of this article is to investigate whether the isolation, the primacy, and the recency effects in free recall data can be accounted for by experience-dependent synaptic plasticity in neural cells. In particular, it is shown that all these effects can be accounted for by dynamically changing the threshold for increasing or decreasing synaptic strength, depending on the postsynaptic activity in a neural network, whereas none of these phenomena are obtained if the threshold is constant.

The isolation, primacy, and recency effects also share properties with what is known about synaptic plasticity, namely, dependency of prior events. Strengthening or weakening of synaptic plasticity depends on this level of activity, as well as the history of activity. High postsynaptic activity leads to strengthening, or long-term potentiation (LTP), and low activity leads to weakening, or long-term depression (LTD) of synaptic strengths (Lynch, 2003). The threshold between strengthening and weakening is set by the recent average activity, so that a high activity leads to a higher threshold or more LTD, and low activity to a low threshold and more LTP. These properties enhance stimulus-selective neural responses when exposed to a range of stimuli in artificial neural networks (Bienenstock, Cooper, & Munro, 1982; Clothiaux, Bear, & Cooper, 1991). These neuropsychological findings are of great relevance to the isolation effect. An isolated stimulus is represented by very different cells than those representing a nonisolated stimulus and is more likely to be recalled because recent activity has been low in these cells, leading to greater strengthening or LTP, and because it is less susceptible to weakening or interference from subsequent nonisolated items.

The mechanism for dynamically changing the threshold also has a computational advantage. That is, it addresses a fundamental problem in neural networks, namely, the difficulty of encoding and retrieving correlated patterns (Lewandowsky, 1991; Lewandowsky & Li, 1995; Metcalfe, 1993a; Simon & Lewandowsky, 2002)—for example, a list of contextually related items that are encoded in the isolation effect. It is shown how this problem can be dealt with by dynamically changing the LTP–LTD threshold.

The article is organized as follows. First, behavioral data on the isolation, the primacy, and the recency effects are reviewed. Second, neurophysiologic findings on synaptic plasticity, in particular LTP–LTD and bounded synaptic plasticity, are reviewed. Third, a model based on these neurophysiologic findings is suggested to account for the behavioral data. Fourth, the key mechanisms of the model are explained, and the model is fitted to empirical data, followed by a general discussion.

1.1. The isolation effect

The isolation is the robust and strong effect where an isolated item is better remembered in a free recall test than an item that is consistent with the surrounding context. The isolated item can be made different from the surrounding context in a variety of ways (for a review, see Schmidt, 1991): by a physical difference, for example, one word printed in red and the other words printed in black (Wallace, 1965); by a conceptual difference, for example, the word *car* embedded in a list of various flower names (Hunt & Mitchell, 1982); by showing nude pictures in a series of more typical magazine pictures (Ellis, Detterman, Runcie, McCarver, & Craig, 1971); by consistency or inconsistency with the subjects' schema, for example, whether an ashtray is seen either in a preschool classroom or in an office (Pezdek, Whetstone, Reynolds, Askari, & Dougherty, 1989); by predictability, for example, a single word that was either predictable or not predictable from the a preceding text (O'Brian & Myers, 1985); by words rated as more distinctive than words rated as not distinctive (Hunt & Elliot, 1980); by bizarre versus nonbizarre imagery (Einstein, McDaniel, & Lackey, 1989); and so on, where the isolated, deviating, or different condition is better recalled or remembered than the corresponding nonisolated condition.

A traditional view of the isolation effect is to emphasize salience, distinctiveness, or interference; however, as argued here, the isolation effect cannot easily be explained by these concepts. For example, Jenkins and Postman (1948) were first to propose that differential attention is the mechanism mediating the effects of distinctiveness. Green (1956) argued that the isolation effect results from a surprise induced by the change from preceding items. Murdock (1960) defined *distinctiveness* as the extent to which a given stimulus "stands out" from other stimuli and noted that the concept of distinctiveness refers to the relation between a given stimulus and one or more other stimuli to compare with. That is, if there are no comparison stimuli, the concept of distinctiveness is simply not applicable.

Although salience and distinctiveness are intuitively appealing as theoretical constructs for the isolation effect, there are problems with this view. Already in the original article introducing the isolation effect, von Restorff (1933) argued against the necessity of salience at encoding for superior memory performance. Unlike the standard practice of placing the isolated item in the middle of the list, von Restorff placed this item at the second position. At the presentation of this item, no context has been established and participants should therefore not conceive of this item as being salient. Yet, an isolation effect was observed. Later research has also found the isolation effect when the isolated item is placed at the first serial position (Kelley & Nairne, 2001; McConnell, Sherman, & Hamilton, 1994; Pillsbury & Rausch, 1943).

Von Restorff (1933) suggested that the isolation effect occurs due to impoverished memory of the nonisolated items rather than the salience of the isolated item. That is, it was argued that the isolation effect occurs because of interference among the similarity of the nonisolated items, making them less easily retrievable than the isolated item. This interference account makes different predictions depending on whether the comparison of the isolated item is made to an item at the same serial position in control lists of homogenous or heterogeneous items. It suggests that the isolated item should be better retrieved than a corresponding item in a homogenous control list consisting of items from the same category, where the interference is large; this has been found experimentally (von Restorff, 1933). The interference account also pre-

dicts that the isolated item should be recalled equally well compared to a corresponding item in a heterogeneous control list consisting of categorically different items, when a heterogeneous control list consists of widely different items,—for example, a number, a letter, a word, a color, a figure, and so on. This is predicted because the similarity between the isolated and the nonisolated item in the experimental condition is equally different to the similarity between items in the heterogeneous control condition, so the performance also should be equal. Contrary to this prediction, experimental data show that isolated items in a context are better retrieved than a corresponding item in a heterogeneous control list. However, von Restorff (1933, p. 318) argued that this finding can be accounted for, in line with the Gestalt tradition, because the heterogeneous case “will be more difficult to memorize than the isolated case even though the items are objectively different from one another, because the *continuous* great difference between items creates the impression of a uniform whole” (emphasis mine).

Hunt and Lamb (2001) emphasized that the processing of similarities and differences are important for understanding the isolation effect. The similarities are related to organizational processing that provides retrieval support by delineating a search set. However, organizational processes are insufficient for memory performance because all members share the same organizational cues. Distinctive processes contribute to item-specific information. From this perspective nonisolated items in an isolation context are experienced under less than optimal circumstances.

Hunt and Lamb (2001) tested this analysis by encouraging distinctive processing of background items and homogenous control list items via a difference judgment-orienting task. They argued that such processing would draw attention to differences among the items and should add the necessary distinctive processes to the ongoing organizational processing. The data confirmed this hypothesis, and the isolated item was no better remembered than background items or the homogenous control list item after the difference judgment-orienting tasks.

Few quantitative models of the isolation effect exist. However, Nairne (1990) proposed a feature model that was briefly applied to the isolation effect by Kelley and Nairne (2001). In this model traces that have not decayed from primary memory are used as cues to sample recall candidates from secondary memory. At test, traces are refreshed by comparing them with possible candidates in a secondary memory search set. If a particular feature is unique, then performance for the item improves. Neath (1993) suggested a distinction between processing of context as reflecting the beneficial effect of similarity and distinctiveness as reflecting the beneficial effect of dissimilarity. This context-distinctiveness view was applied to the isolation, the primacy, and the recency effect over different time scales. To summarize, the isolation effect is a nontrivial phenomenon. It cannot easily be accounted for by mechanisms such as salience, distinctiveness, or by interference.

1.2. *The primacy effect*

The primacy effect in free recall is the phenomenon that the first few items on a list are better recalled than items in the middle of a list. The dominating account of the primacy effect in free recall is repetition in short-term memory (Atkinson & Shiffrin, 1968; Phillips, Shiffrin, & Atkinson, 1967). That is, items at the beginning of the list are rehearsed more because there are

fewer other items that interfere with the repetition compared to items later in the list. However, this account seems unsatisfactory as a complete account for all aspects of the primacy effect. Attempts made to eliminate rehearsal, for example—as reviewed later—by a short presentation time, do not eliminate the primacy effect. However, it is frequently claimed that although the primacy effect is not eliminated, the conditions where rehearsal is eliminated have a lower primacy effect than conditions where rehearsal is possible.

However, this conclusion is confounded by the fact that elimination of rehearsal, for example, by using a fast presentation rate, not only attenuates the performance of the first items in the list but also attenuates the overall performance level. To avoid this confound it would be useful to have a quantitative measurement of the primacy effect that takes this confounding into account. To the author's knowledge no quantitative measurement of the size of the primacy effect has been introduced in the literature. Therefore *the primacy ratio* (PR) measurement of the primacy effect is introduced here. This measurement does not depend on the overall performance level. It is defined as one minus the ratio of the probability of correct recall of the last serial position in the primacy effect and the probability of correct recall of the first serial position.¹ The last serial position in the primacy effect is defined as the first serial position followed by a serial position with a higher performance.² The larger the value of the PR, the stronger the primacy effect is. A zero value indicates no primacy effect, whereas a positive value indicates a primacy effect. The maximum possible PR is 1.

Notice that the PR may provide a different perception of what constitutes a primacy effect. For example, if the probability of recall at immediate testing for serial position 1 to 5 is 1, 0.5, 0.3, 0.1, and 1, then most research would conclude that there is a primacy effect due to the large slope at the early positions. However, assume that after a considerable delay the performance has declined with a factor 1/100 to 0.010, 0.005, 0.003, 0.001, and 0.010, then many researchers would conclude that there is no primacy effect because the slope is approximately flat. This indicates that measuring the primacy by the slope is problematic, given that one is interested in inferring conclusions of what is happening during encoding. That is because this measurement is very dependent on other strength factors, for example, the retention interval. However, the PR measurement is less sensitive to similar strength factors, and in this example the PRs for the two conditions are actually identical $1 - 0.1/1 = 1 - 0.001/0.010 = 0.9$.

Using the PR measure, available data show a strong PR for a large number of manipulations aimed to eliminate rehearsal. The number in parentheses shows the PR. Primacy effects are found in the following conditions: fast study time; a positive PR is found if items are presented so fast that rehearsal is impossible or unlikely. For example, the PR in Wixted and McDowell (1989) was larger (PR = 0.50) under fast conditions (0.50 sec per item) than under slow conditions (PR = 0.41; 2.5 sec per item). All participants in the fast condition reported that they did not rehearse, suggesting either that the primacy effect is independent of rehearsal, or that rehearsal occurs without awareness (for similar findings, see Glanzer & Cunitz, 1966). A primacy effect was found following incidental instructions when participants have no reason whatsoever to rehearse because they are simply not expected to be tested (PR = 0.23 in Darley & Glass, 1975; PR = 0.67 in Gershberg & Shimamura, 1994) and for stimuli materials that are difficult to rehearse—for example, pictures (PR = 0.52 in Watkins & Peynircioglu, 1983). Primacy is found during rehearsal suppression, or fixed rehearsal, where participants are instructed to only rehearse the presently encoded item (PR = 0.43 in Fischler, Rundus, &

Atkinson, 1970); it is also found when participants are allowed to overtly rehearse any item (PR = 0.56 in Fischler et al., 1970). A pronounced primacy effect also occurs under continuous distractor tasks—for example, concurrent counting backward (PR = 0.60; Modigliani & Hedges, 1987); it also occurs when rehearsal is encouraged (PR = 0.39 Modigliani & Hedges, 1987). The primacy effect is found in other species besides humans, where elaborated strategies such as rehearsal are less likely to occur—for example, in primates and pigeons (PR = 0.20 in Wright, Santiago, Sands, Kendrick, & Cook, 1985). Due to the arguments and data presented previously, it is reasonable to conclude that factors other than rehearsal contribute to the primacy effect.

Notice also that this discussion has focused on the problem with rehearsal as an account for the primacy effect. It is of course possible for items to be maintained for a limited time in short-term memory without a rehearsal and that the time in short-term memory is related to the primacy effect; however, it is not immediately obvious how such a framework accounts for the improved performance for isolated items without evoking additional mechanisms such as rehearsal or distinctiveness.

1.3. *The recency effect*

The recency effect in free recall is the phenomenon that the most recently encoded items are, on immediate retrieval, better recalled than items in the middle of the list (e.g., Murdock, 1962; Tan & Ward, 2000). The original explanation of the recency effect was that items first reside in a capacity-limited short-term buffer from which the items are immediately retrieved. Subsequent retrieval of earlier items are recalled from a more permanent long-term memory store (Atkinson et al., 1968). This dual-store model was later challenged by empirical data showing a “recency” effect also for retrieval from long-term memory. In particular, recency was found in a continuous distractor task, where each encoded item is followed by a distractor task of different lengths. A recency was found when the length of the distractor for the last item was not longer than the length of the earlier distractors (Bjork & Whitten, 1974; Glenberg, Bradley, Kraus, & Renzaglia, 1983). However, several lines of evidence have suggested that long-term recency (LTR) is different from short-term recency (STR). For example, a negative recency effect is found in STR but not in LTR. This effect occurs when participants first make an immediate recall from the lists. Following a delay, a final free recall test is given where the performance on the last few items are lower than items in the middle of the list (Craig, 1970). STR is sensitive to output order, whereas LTR is not (Dalezman, 1976). STR is insensitive to damage to the medial temporal lobe (Carlesimo, Marfia, Loasses, & Caltagirone, 1996). Kahana (1996) showed that items that appear nearby during encoding are more likely to be recalled together than items that are far apart during encoding. This effect is asymmetrical so that forward transitions are more likely than backward transitions. However, this effect is much more pronounced during recency items compared to items from the middle of the list.

Several variables affect prerecency items but not recency items—for example, list length, word frequency, semantic similarity, and proactive interference in free recall (Craig & Birtwistle, 1971; Glanzer, 1972). Finally, several variables that influence short-term memory or the recency effect—for example, visual or auditory modality—do not influence the primacy effect (Craig, 1970). Later theories have emphasized the role of a continuously changing con-

text, and they account for the recency effect by residual overlap between the retrieval context and the encoding context (e.g., Davelaar, Goshen-Gottstein, Ashkenazi, Haarmann, & Usher, 2005; Howard & Kahana, 2002). Here it is shown how primacy, recency, and isolation effects are predicted by an adaptive LTP–LTD threshold and bounded synaptic strength.

1.4. Adaptive LTP–LTD threshold and synaptic plasticity

A fundamental issue in memory is how synaptic plasticity occurs due to experience. A number of theoretical analyses of plasticity in the visual cortex (Bienenstock et al., 1982; Miller, 1990; von der Malsburg, 1973) are based on the Hebbian postulate, which states that synapses become more effective when the presynaptic activity consistently correlates with the postsynaptic response beyond some threshold (Hebb, 1949/2002). In addition, weakening occurs when pre- and postsynaptic activities are consistently uncorrelated (Frégnac & Shulz, 1994).

Of particular interest, Bienenstock and colleagues (Bienenstock et al., 1982; Cooper, 1995; Cooper, Liberman, & Oja, 1979) suggested that active synapses are potentiated when the total postsynaptic response exceeds a critical value called the *modification or adaptive threshold*. Furthermore, synapses are depressed when the total postsynaptic response is less than the adaptive threshold but larger than zero.

By stipulating that the adaptive threshold varies as a function of the average postsynaptic response, Bienenstock et al. (1982) showed that neuronal responses would evolve to stable selective states. Later computational studies have shown that this Bienenstock–Cooper–Munro theory can account for a wide variety of experimentally observed visual-deprivation effects on cortical receptive fields (Clothiaux et al., 1991). Furthermore, the key assumptions of the Bienenstock–Cooper–Munro theory have a physiological basis in the visual cortex (Bear, 1996).

Bear, Cooper, & Ebner (1987) proposed that the adaptive threshold corresponds to a critical level of Ca^{2+} entry through the postsynaptic NMDA receptor channel (NMDA receptors play an important role in synaptic plasticity). The influx of Ca^{2+} in the postsynaptic cell triggers a cascade of processes leading up to LTP. This proposal was based on evidence in the CA1 region of hippocampus where NMDA-receptor-mediated Ca^{2+} entry can trigger synaptic LTP (Nicoll, Kauer, & Malenka, 1992).

Thus, these studies suggest that the adaptive threshold for synaptic potentiation and depression varies as a function of the history of the postsynaptic activity. After a period of cortical inactivity of the postsynaptic cell, the adaptive threshold should be reduced. Following a period of high activity of the postsynaptic cell, the adaptive threshold should be increased. Furthermore, activity in the presynaptic cell is required for both the increase and the decrease in synaptic strength.

1.5. Bounded synaptic plasticity

Synaptic plasticity depends on experience. However, even if LTP protocols are repeated for a large number of trials, synaptic strength does not grow unlimited. Instead, repeated inductions of LTP and LTD protocols in rat motor cortex show that synapses have a certain modifica-

tion range. That is, if the protocol known to induce LTP (or LTD) is repeated, then over time the LTP (or the LTD) reaches an asymptote or saturates, and the modification range is the distance between the saturated LTD and the saturated LTP. The synaptic modification range has been studied during motor skill learning in rats, where LTD saturation has been found to occur approximately 60% below baseline, and LTP 160% above baseline (Riout-Pedotti, Friedman, & Donoghue, 2000). The exact mechanism for the synaptic modification range is not completely understood; however, one limiting factor of LTP may be that the number of synaptic vesicles is limited, and the number of transmitter molecules released is bounded to packets of about 6,000 (for reviews, see Ceccarelli & Hurlbut, 1980; Jahr, 2003; Katz, 1969; Steinbach & Stevens, 1976). Sikström (2002) implemented a synaptic modification range in an artificial neural network by introducing synaptic boundaries. Synaptic boundaries cause interference so that the expected strength attenuates according to an exponential function for single weights as new items are encoded. Introducing variability in the learning rates causes the expected strength of the aggregated synaptic strengths to show a power-function forgetting curve, characterized by decreasing rate of forgetting during the retention interval. To produce exact power-function forgetting curves, the variability in the learning rate must be generated from a power function (for a mathematical proof, see Sikström, 2002).

1.6. Computational problems with encoding correlated patterns using Hebbian learning

The isolation effect occurs when a different item is presented in a context. The context consists of items that are similar in some dimensions. In distributed neural systems, similarity may be measured by the mathematical concept of correlation between representations of distributed patterns. Encoding correlated patterns is problematic in some neural networks because the overlap in representation between correlated patterns tends to destroy the memory capacity (Lewandowsky, 1991; Lewandowsky & Li, 1995; McCloskey & Cohen, 1989). Hebbian learning is very efficient for uncorrelated patterns; however, it quickly breaks down for correlated patterns. Encoding correlated patterns tends to increase the strength not only of the encoded pattern, but also of the previously encoded patterns that are correlated with the currently encoded pattern (to a degree that depends on the correlation). In the context of the isolation effect this means that Hebbian learning increases the strength of the nonisolated items (correlated patterns) more than the strength of the isolated item (uncorrelated pattern). Thus, Hebbian learning makes the opposite prediction from what is found in the isolation effect. Furthermore, strengthening of correlated patterns in Hebbian learning tends to make the network unable to retrieve correctly. (For related issues concerning learning of correlated patterns, see McClelland, McNaughton, & O'Reilly, 1995; Minsky & Papert, 1969/1988; Rumelhart, Hinton, & Williams, 1986.)

There are different ways to deal with the problem of encoding correlated patterns. For example, Hebbian learning can be replaced by error-correcting learning rules; however, these learning rules have been argued to be less plausible mechanisms in biological neural networks (Crick, 1989; Zipser & Andersen, 1988; for a different view see O'Reilly, 1996). Here it is focused on how to modify Hebbian learning to deal with correlated patterns by using an adaptive LTP–LTD threshold.

1.7. *Justifications*

Previous theories of the isolation effect have focused on theoretical concepts such as salience, distinctiveness, and interference. Earlier accounts of the primacy effect have emphasized rehearsal, and the recency effect is typically accounted for by a short-term buffer. The purpose of this article is to show that an adaptive LTP–LTD threshold accounts for the isolation in the middle of the list, the primacy, and the recency effects. Furthermore, when the adaptive threshold is combined with synaptic boundaries, an isolation effect is found at the first serial position, in comparison with a heterogeneous control list, and it also accounts for power-function forgetting curves. In particular the memories for isolated associations (MIA) model is proposed, which uses a Hebbian learning rule with an adaptive LTP–LTD threshold and bounded synaptic strengths (MIA is a subset of a more general model for episodic memory; see Sikström, 2000, 2001, 2002; Sikström & Jaber, 2002). The proposed learning rule is also useful from a computational perspective because the neural network performance is also efficient in an environment where the input patterns are correlated.

MIA simultaneously explains the isolation effect from three different levels, namely, neural, computational, and cognitive–behavioral levels. At the neural level MIA is based on current theories of synaptic plasticity or on Hebbian learning with an adaptive LTP–LTD threshold and bounded synaptic strength. At the computational level it is shown that Hebbian learning cannot efficiently encode correlated patterns; however, this is possible when Hebbian learning is combined with an adaptive threshold. At the behavioral level it is shown that Hebbian learning with adaptive threshold accounts for the isolation effect and serial positions effects.

Memory can be disturbed by interference. Some models emphasize interference at retrieval (Gordon, Hendrick, & Levine, 2002), whereas others predict interference at storage (Lewis, 1996). MIA makes the biologically plausible assumption that interference occurs at storage, whereas no interference mechanisms at retrieval are needed to account for the relevant phenomena. Furthermore, this storage interference may be either proactive, which is interference occurring prior to encoding, or retroactive, which is interference occurring after encoding of the to-be retrieved item.

The predictions of the model can be summarized by retroactive and proactive interference modified by the LTP–LTD threshold and synaptic boundaries as follows: The primacy effect is accounted for because the adaptive LTP–LTD threshold is low during encoding of the first few items; this lack of proactive interference leads to synaptic strengthening during encoding (LTP) and better performance compared to items in the middle of the list. The adaptive threshold is high following encoding of the last list item, and the recency effect occurs because no retroactive interference or LTD has yet occurred on the last item or items at immediate testing. The isolation effect in the middle of the list occurs because the adaptive threshold is low for the isolate, and this lack of proactive interference increases the synaptic strength (LTP). Furthermore, the presynaptic cells tend to be inactivated following the isolated item, causing a blocking of LTD, or a lack of retroactive interference, which also leads to an isolation effect at the first serial position. Synaptic boundaries cause weaker retroactive interference on an isolated item presented in a congruent list than on a corresponding item in a heterogeneous list, yielding a better performance for the isolate compared to a heterogeneous control list. Synaptic boundaries also account for power-function forgetting curves due to retroactive interference.

This article is organized in the following way. The MIA model is first presented verbally, then implemented, and simulated in a neural network (see the Appendix for an explicit solution and technical details). The effect of LTP, LTD, and interference on performance is specified. The mechanisms accounting for primacy, recency, and isolation effects are discussed, and the model is fitted to empirical data. The necessity of adaptive thresholds and synaptic boundaries are emphasized by showing that these mechanisms are critical to account for the behavioral data. A discussion follows, including comparisons to other models.

2. The MIA model

A longstanding goal in the investigation of behavior is to understand how the cognitive–psychological level and the neurological level can be linked. MIA makes several connections between these two levels. For example, the cognitive concept of memory representation is assumed to be mapped into the distributed pattern of activity in a large number of neural cells. The cognitive representation of a categorized list is assumed to be mapped into the activity of a subset of neural cells so that, for example, the representation of words belonging to the category “fruit” activates a subset of neural cells to a higher degree than words belonging to other categories. The cognitive concept of an item’s distinctiveness maps to that the pattern of activation, that is the distinctive item is largely uncorrelated, or different from, the representation of the nondistinctive items. The cognitive concept of encoding is mapped onto the neurological finding of experience-dependent synaptic plasticity. In this sense MIA combines the cognitive psychology and neural levels with an emphasis to explain the phenomena found at the cognitive–behavioral level.

Hebbian learning is based on the idea that the synaptic strength between two cells that are simultaneously active should be increased. This encoding principle provides an efficient rule for changing synaptic strengths for uncorrelated patterns. For example, if the presynaptic cell is active at retrieval, then the increased synaptic strength will tend to also activate the postsynaptic cell. In this way the presynaptic cell reactivates the encoded state of the postsynaptic cell. Whereas if the presynaptic cell is inactive at retrieval, then the increase in synaptic strength is of no use, and the postsynaptic cell will tend to be inactivated.

However, there is more to synaptic plasticity than just increasing the synaptic strength. At some point the synaptic strength must also be decreased, otherwise all synapses will increase over time, and all postsynaptic cells would be activated with minimal input. Thus, there must be some mechanism that decreases the synaptic strength. In Hebbian learning the synaptic strength is decreased if the postsynaptic cell fires with a low frequency (moderately inactive),³ and the presynaptic cell is active.

The following notation is used. Assume that the postsynaptic cells are either in an active state represented as 1, with a probability of p , or an inactive state represented as 0, with a probability of $1 - p$. Let the parameter $0 \leq a \leq 1$ represent the adaptive LTP–LTD threshold, so that the increase in synaptic strength for an active postsynaptic cell is $1 - a$, and the decrease in synaptic strength for an inactive postsynaptic cell is $0 - a$ (i.e., the change in synaptic strength equals the activity minus the adaptive threshold). Furthermore, synaptic changes only occur if the presynaptic cell is active (Bienenstock et al., 1982; Cooper, 1995; Cooper et al., 1979).

To make the network usable after repetitive learning of items, the increase and decrease in synaptic strength should balance so the sum of changes over all synapses in the network is zero. To accomplish this equilibrium, the sum of the expected increase and decrease of synaptic strength must be zero.⁴ The expected increase equals the probability that the postsynaptic cells are active (p) times the increase in synaptic strength ($1 - a$). The expected decrease equals the probability that the postsynaptic cell is inactive ($1 - p$) times the decrease in synaptic strength ($-a$). Thus, the sum of the expected increase in synaptic strength and the expected decrease in synaptic strength is zero, given that $a = p$; that is, $p(1 - a) + (1 - p)(-a) = p - a$. This shows that if the probability that the postsynaptic cell activity (p) is large, then the adaptive threshold must also be large (a). If the postsynaptic activity is low, then the adaptive threshold must also be low.

To keep the expected synaptic strength at zero, changes that occur with a low probability must be associated with a large magnitude of synaptic strength. Conversely, changes that occur with a high probability must be associated with a low magnitude of synaptic strength. This implies that a postsynaptic cell that is frequently inactive should make large increases in the synaptic strength on those few occasions when it is active and small decreases when it is inactive. In contrast, a postsynaptic cell that is frequently active should make small increases in the synaptic strength when it is active and large decreases when it is inactive. For example, if cells are active 2 times out of 10 ($p = a = 2/10$), then the increase in synaptic strength ($1 - a = 0.8$) during the two active states is $0.8 + 0.8 = 1.6$, and the decrease in synaptic activity ($-a = 0.2$) during the eight inactive states is $-0.2 - 0.2 - 0.2 - 0.2 - 0.2 - 0.2 - 0.2 - 0.2 = -1.6$. If the probability that the postsynaptic cell is active is larger than the adaptive threshold, then the synaptic strength will grow unbounded over time, whereas if this probability is less than the adaptive threshold, then the synaptic strength will decrease unbounded over time.

At retrieval an encoded pattern is reinstated as a retrieval cue (see later details). The synaptic strengths from connected active cells are added to form a net input to a cell. That is, the net input to a cell equals the sum of synaptic strengths from all active cells connected with this cell. The synaptic strengths from inactive cells do not contribute to the net input. In the MIA model the performance is based on the net input.

2.1. Implementing MIA in a connectionist neural network

A verbal description of MIA was presented previously. In the Appendix, MIA is formalized in a neural network model. The Appendix also provides an explicit analytical solution to the model when the performance is measured by net inputs. This solution is used for fitting the model with empirical data. The model is also investigated by simulations that allow for testing other performance measurements. By convention in the literature, the synaptic strength is referred to as weights and cells as nodes in the neural network implementation of MIA. Readers who are less interested in technical details of neural network models may skip the Appendix without loss in comprehension of the model. That is, the mechanisms described previously are also valid for the formal model presented in the Appendix.

MIA is identical to a Hopfield neural network model (Hopfield, 1982, 1984) with the following two modifications. First, the Hopfield network uses a fixed threshold, whereas MIA uses a threshold that adapts according to the history of activity of the corresponding

postsynaptic node. In the presynaptic node the adaptive threshold is zero, so weight changes only occur when the presynaptic cell is active. Second, the Hopfield network uses unbounded weights with a constant learning rate, whereas MIA uses bounded weights with variability in learning rates (Sikström, 1999, 2002).

2.1.1. Representation

The network consists of one layer that is fully interconnected (Fig. 1). Episodic memory is typically represented as associations between items and encoding contexts (Sikström, 2001). However, here we only represent items, and do not represent the encoding context because this context is not necessary for our simulation purposes (and excluding the encoding context keeps the implementation as simple as possible). Items are represented in distributed activity patterns in nodes. Patterns are generated as follows. Heterogeneous lists (uncorrelated lists, left panel of Fig. 1) are generated by setting nodes to an active state (1) with a probability of p , and otherwise to an inactive state (0). There is a total of N number of nodes. Homogenous list (correlated patterns, right panel of Fig. 1) are generated by activating a subset, consisting of N_c ($< N$) number, of nodes with a probability of p_c . The remaining $N - N_c$ nodes are inactive. The

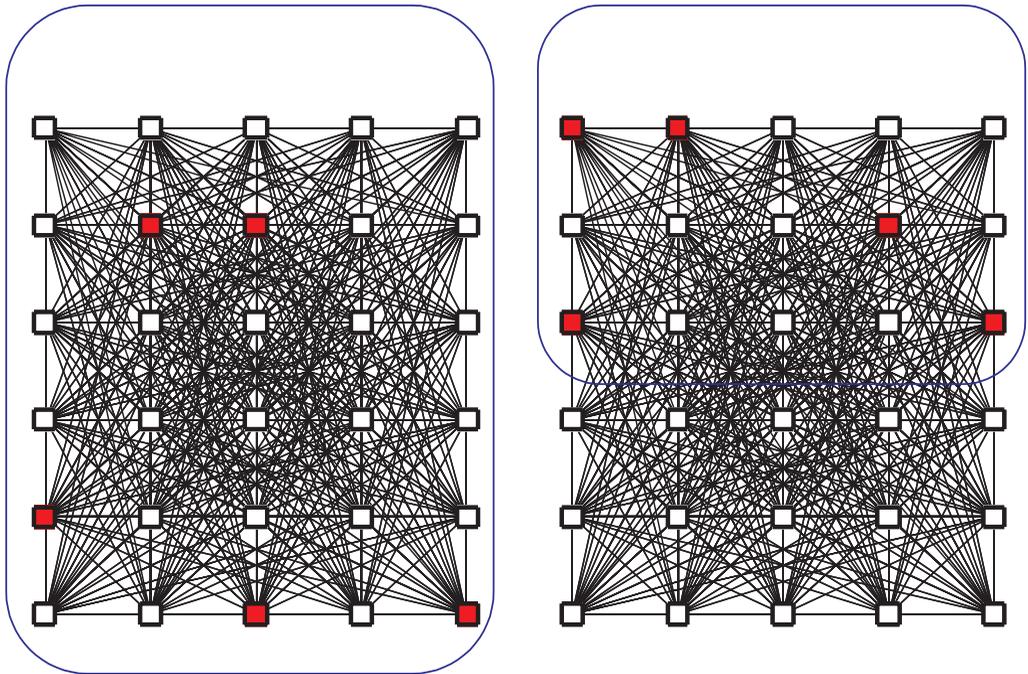


Fig. 1. The topology of the MIA network.

Note. The squares represent active (red) or inactive (white) cells, and the lines are the connections (i.e., synapses) between cells. Heterogeneous lists and isolated items are generated by randomly activating cells from the whole network (left network), whereas congruent lists are generated by activating a subset of the cells (e.g., upper half) in the network (right network).

expected number of active nodes is the same for any pattern. That is the expected number of active nodes is the same independently of whether the pattern is taken from an uncorrelated or a correlated category. From this it follows that a node within the correlated category is more likely to be active than a randomly chosen node (i.e., $p < p_c = N/N_c p$). This will also make the adaptive threshold larger for nodes within a category than for nodes outside a category.

2.1.2. Conditions

There are three conditions. (a) In the heterogeneous condition all patterns are generated by randomly activating nodes with a probability of $p = 5/30$ (left panel of Fig. 1). (b) In the Position 9 condition an isolated pattern is presented at Serial position 9, which is generated in the same way as patterns in the heterogeneous condition. The other nonisolated patterns are generated from a subset of the nodes $N_c = 15 < N$ so that the probability that these nodes are active is $p_c = p N/N_c = 5/15$ (right panel of Fig. 1). (c) The Position 1 condition is identical to the Position 9 condition, except that the isolated pattern is presented at the first serial position. The expected number of active nodes is 5 in all conditions. The list length in experimental free recall data typically is 15 to 20 items (e.g., 15 in Wixted & McDowell, 1989; 20 in Tan & Ward, 2000), and here 18 patterns are used in all conditions.

2.1.3. Encoding

The adaptive threshold is implemented in the postsynaptic cell by calculating a time weighted average of activity, where more recent activities are weighted heavier than older ones (see the Appendix for details). A constant, $0 \leq c < 1$, determines the time span of the weighted average, where $c = 0$ sets the adaptive threshold to the activity of the previous encoded pattern, and $c > 0$ weights the average over several previously encoded patterns. In this article $c = 0$; however, all predictions of the model also hold for $0 < c < 1$, where larger values on c tend to produce a flatter primacy effect. Furthermore, weight changes only occur when the presynaptic cell is active.

2.1.4. Retrieval

Retrieval starts with cuing the network by reinstating one of the encoded patterns (i.e., corresponding to the to-be-retrieved serial position). Technically, this is equivalent to a cued recall test (where the cue and the target pattern are identical).⁵ Following cuing, the performance was calculated in a simplified way and in a slightly more complex way. The simplified performance measure is simply the sum of the net inputs for nodes encoded in the active state, where all nodes are clamped to the input pattern. This measurement was used due to its simplicity and because it resembles a feedforward network. Furthermore, the model was also simulated with relaxation (i.e., where the retrieved activity is iteratively used as new cues in the network), which is the retrieval procedure typically used in a standard recurrent Hopfield network. The encoded pattern of activation was reinstated. The net input to a randomly chosen node was calculated, and the node was set to an active state if the net inputs exceeded 0.86 (this value was chosen to obtain a good overall performance); otherwise it was set to an inactive state. This was repeated until there was no change in activity during 60 consecutive iterations, or until a maximum of 150 iterations. An item was considered to be successfully recalled if the recalled activity pattern exactly matched the encoded activity pattern (i.e., the cue pattern); otherwise it was scored as unsuccessfully recalled.

2.1.5. Initialization

At time zero, prior to encoding of each list, the following initializations were made. The learning rates and the initial weights were initialized as described in the Appendix (i.e., the learning rates were for each weight initialized to a random value taken from a certain distribution, and the weights were initialized to a uniformly random value between the lower and the upper boundaries). The modification thresholds were initialized to zero, corresponding to the assumption that participants have not encoded anything just prior to the list learning.

2.1.6. Simulations

After initializations, 18 patterns were encoded, starting with pattern with Serial position 1, 2, 3, to 18. Following encoding of the last pattern, the performances of the encoded patterns as a function of serial position during encoding (i.e., from 1 to 18) were calculated. Furthermore, the performances for two “distractor” items (not encoded) were calculated. The item labeled “U” is an isolated, or uncorrelated, distractor item that has not been encoded, and the item labeled “C” is a nonisolated, or correlated, distractor item that has not been encoded. The results presented in the figures are averaged over 500 list simulations. Each list simulation was different because the weights and the encoded patterns were reinitialized to different values.

2.1.7. Parameters

The number of parameters in the MIA model is kept at a minimum. In the simulation the parameters were set as follows (see the Appendix for a detailed specification of the parameters). There are three parameters describing the patterns, and they influence how similar or dissimilar the nonisolated items are in relation to the nonisolated item and the heterogeneous list: the total number of nodes $N = 30$, the number of nodes in the correlated patterns $N_c = 15$, and the probability of activating nodes $p = 5/30$ in the whole network (corresponding to $p_c = 5/15$ in the correlated patterns).

There are three parameters describing the distribution of learning rates, and they determine the shape of the forgetting curve: $x_{max} = 3$ (this value determines the smallest learning rate in the distribution of learning rates. The higher the value, the longer the retention interval will show a power-function forgetting curve), the boundary $b = 1$ (which is a scaling variable), and $\alpha = 0.5$ determines slope (i.e., the exponent in the power-function) of the forgetting curve. This value was set so that it corresponds to the slope that is typically found in empirical data (see, for example, Rubin, Hinton, & Wenzel, 1999; Rubin & Wenzel, 1996). Finally, there is one parameter describing the rate of adaptation in the LTP–LTD threshold $c = 0$, which for example, influences how many serial positions will be influenced by the primacy effect.

3. Results

3.1. Performance is influenced by the adaptive threshold (LTP–LTD) and synaptic boundaries

The synaptic plasticity and the adaptive threshold during encoding, or after encoding, influence performance in two different ways (i.e., in MIA, LTP and LTD have the following particu-

lar properties): LTP is the increase in synaptic strength that occurs when the presynaptic activity is active and the postsynaptic cell is activated during encoding of an item and results in an increased probability of retrieving that item. LTD is the decrease in synaptic strength that occurs when the presynaptic activity is active and the postsynaptic cell is inactivated after encoding of the to-be-retrieved item.

3.1.1. LTP

In MIA increase in synaptic strength or LTP (Fig. 2c) occurs during encoding of an item, given that all of the following conditions are fulfilled. LTP occurs if the (a) presynaptic cell is active (Fig. 2a) and (b) if the adaptive threshold is lower than the postsynaptic activity. The latter occurs when the postsynaptic cell is active (Fig. 2b), and the postsynaptic cell was inactive at the previous time step or steps (Fig. 2b, at the previous time step when the adaptive threshold is averaged over a short time interval $c = 0$, or several time steps for $c > 0$). No LTP occurs if one or more conditions are not fulfilled.

3.1.2. LTD

In MIA a decrease in synaptic strength or LTD (right panel of Fig. 2c) occurs, given that the following conditions are fulfilled: (a) the presynaptic cell is active (right panel of Fig. 2a) and (b) the adaptive threshold is higher than the postsynaptic activity. The latter occurs when the postsynaptic cell is inactive (right panel of Fig. 2b), and the postsynaptic cell was active at the previous time step or steps (the previous time step for $c = 0$, or the weighted average of several time steps for $c > 0$). No LTD occurs if one or more conditions are not fulfilled. Notice that LTD diminishes the performance after the to-be-recalled item has been presented (whereas LTP improves performance during the presentation of the to-be-recalled item). LTD occurs after the postsynaptic cell is deactivated, and it requires that the presynaptic cell is active. This means that LTD diminishes performance of the to-be-recalled item more when it is followed by a similar item.

3.1.3. Synaptic boundaries

The performance is also influenced by synaptic boundaries causing forgetting due to retroactive interference as other items are encoded during the retention interval. For unbounded synaptic strengths (and an appropriately placed adaptation threshold), the encoded synaptic strength is preserved during the retention interval in the expected value of the synaptic strength. Encoding during the retention interval will only add noise (with zero mean) to the synaptic strength, and the amount of noise depends on how often the synaptic strength is updated, whereas the expected value of the strength is intact. However, as synaptic boundaries are introduced, the expected synaptic strength attenuates toward zero during the retention interval. The upper boundary “pushes” the expected strength to smaller values, and the lower boundary pushes it to higher value so that the expected strength approaches zero as the retention interval increases. For a single synapse, associated with a certain learning rate, the expected synaptic strength across the retention interval can be described by an exponential function; however, as different learning rates are combined, the aggregated strength of several synapses tends to approximate a power function.

This retroactive interference is modulated by the adaptive threshold. The retroactive interference increases with the adaptive threshold. A large adaptive threshold causes large interfer-

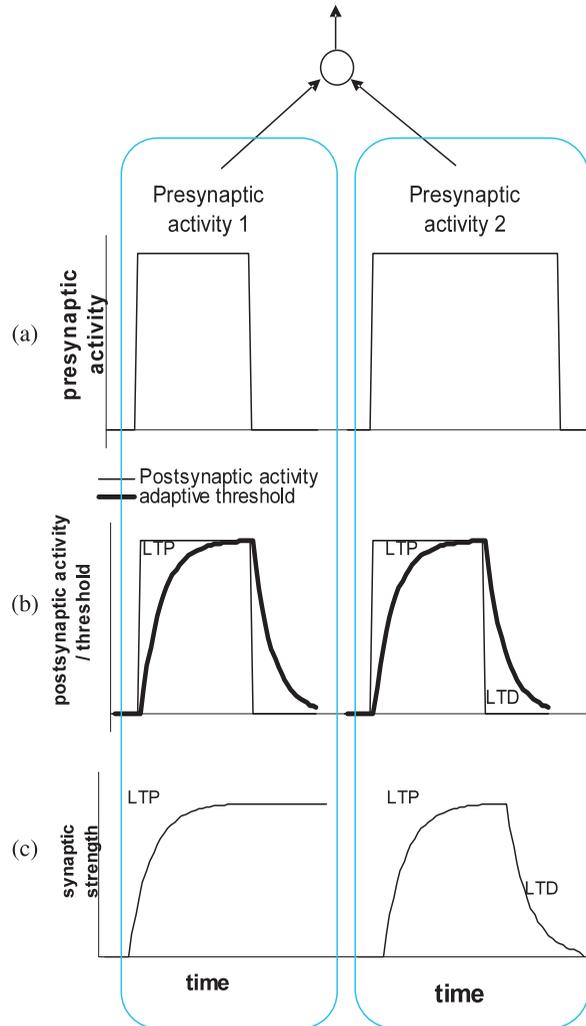


Fig. 2. A Hebbian learning rule with an adaptive LTP-LTD threshold.

Note. (a) The presynaptic activities for two cells. Both cells are activated at the same time. Presynaptic cell 1 (left panel) is deactivated when the postsynaptic cell deactivates, whereas Presynaptic cell 2 (right panel) remains active after the postsynaptic is deactivated. (b) The postsynaptic activity changes from an inactive to an active state and from an active to an inactive state. The adaptive threshold slowly (i.e., depending on the constant c) adapts toward the postsynaptic activity. (c) The synaptic strengths from Presynaptic cells 1 and 2 to the postsynaptic cell increase following pre- and postsynaptic activation to an extent that depends on the adaptation of the adaptive threshold. The synaptic strength reaches an asymptotic value that corresponds to the area labeled LTP. The synaptic strength decreases (LTD) when the postsynaptic cell is inactivated and the Presynaptic cell 2 is active to an extent that depends on the adaptive threshold. The synaptic strength decreases with a value that corresponds to the area labeled LTD. There is no LTD for Presynaptic cell 1 because this cell is inactivated when the postsynaptic cell is inactivated. The figures were produced by keeping the postsynaptic cell active for 35 time units and setting $c = 0.85$ (the extended time scale was chosen to produce smooth lines in the figures).

ence and quick forgetting, and low adaptive threshold causes low interference and slow forgetting. For example, assume that the synaptic strength is 0.5, $a = 0.5$, and the upper–lower synaptic boundaries are ± 0.5 . If the cell is active at the following time step, then the synaptic strength will remain at 0.5 (due to the boundary; if there were no boundary then the synaptic strength would increase by a factor $1 - a$ to 1). However, if the cell is inactive at the following time step, then the synaptic strength will decrease (by a factor $-a$) to 0. Thus, the expected synaptic strength decreases to 0.25 or forgetting occurs because of the bounded synaptic strength. The larger the a , the larger the retroactive interference is in the synaptic strength.

3.2. *MIA accounts for behavioral phenomena by proactive and retroactive interference*

Fig. 3a shows the results from the simulation of MIA using adaptive thresholds and bounded synaptic strength. The results show recency, primacy, and isolation effects. The isolated patterns have a larger net input compared to the heterogeneous control list both at Position 1 and Position 9. How LTP, LTD, and synaptic boundaries account for these phenomena is described in detail later.

3.2.1. *Adaptive threshold and LTP yields a primacy effect due to lack of proactive interference*

At the start of the presentation of a list, the cells have been inactive for a while and no proactive interference exists. Thus, the network starts off with a low adaptive threshold yielding strong encoding and a good performance for the first few patterns in the list. After a few items, the thresholds have adapted and is yielding a relatively weaker encoding and poorer performance compared to the first items. The adaptation of the threshold may be seen as a buildup of proactive interference that attenuates performance across serial positions. Adaptive thresholds yield a primacy effect (Fig. 3a). The PRs are 0.78 (Position 1 condition), 0.38 (Position 9 condition), and 0.15 (heterogeneous condition).

3.2.2. *Adaptive threshold and LTD causes a recency effect due to lack of retroactive interference*

LTD decreases the synaptic strength in synapses that has previously been active due to retroactive interference. Recency items have not yet been influenced by this retroactive interference and are therefore efficiently retrieved. The retroactive interference eventually diminishes the performance, causing the recency effect to disappear on delayed tests.

3.2.3. *Bounded synaptic strength with variability in learning rates yields power-function forgetting curves due to retroactive interference*

In MIA the recency effect can be accounted for by LTD; however, the synaptic boundaries also contribute to forgetting. The boundaries of the synaptic strength make the expected net input attenuate toward zero due to retroactive interference from encoding of other patterns during the retention interval. The rate of attenuation for the expected value of a single synapse is exponential; however, the expected strength of several aggregated synapses produces a power-function forgetting curve, given that there is variability in the learning rates. Synaptic

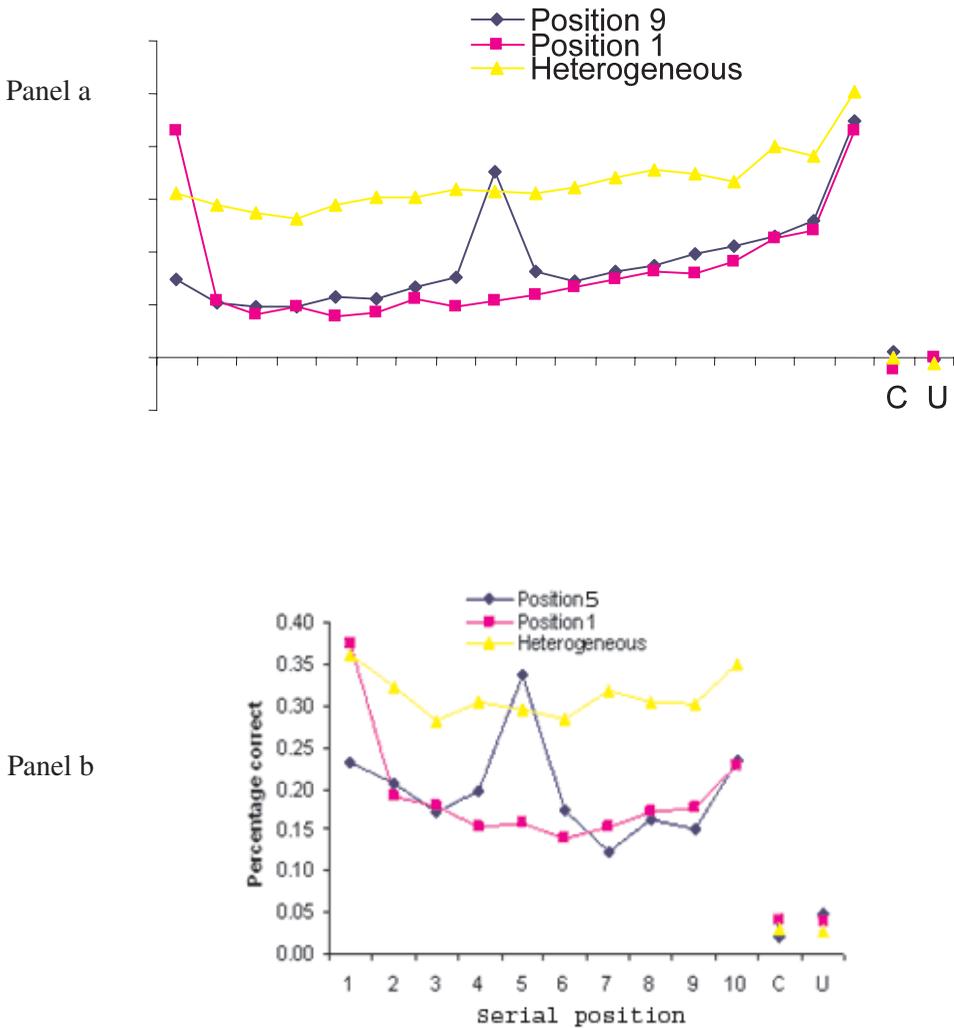


Fig. 3. Predicted serial position curves.

Note. The panels show the prediction of the model as a function of Serial positions 1 to 18. There are three conditions as described in the text. The category labeled “U” shows the averaged results for an isolated item that has not been encoded, and the category labeled “C” shows the results for a nonisolated item that has not been encoded. Panel (a) shows the results as measured by the summed net inputs for cells encoded in the active state, whereas Panel (b) shows the results as measured by the percentage of correctly retrieved patterns. Fewer serial positions are simulated in Panel (b) than Panel (a) to increase the number of correctly recalled items; see the text for details.

boundaries and variability in learning rate yields a forgetting curve that can be described by a power function. A power function is characterized by a diminishing rate of forgetting as the retention interval increases (Fig. 3a). That is, strong forgetting occurs during the recency effect (right part of the serial position curve), whereas forgetting attenuates during the asymptote (middle part).

3.2.4. Adaptive thresholds cause LTP, which improves performance due to release from proactive interference for isolated patterns in the middle of the list

Infrequently active cells will have low adaptive thresholds, and when these cells are activated the LTP will be large. Cells representing the isolated pattern have a lower adaptive threshold than cells representing the nonisolated patterns, leading to better performance for the isolated pattern in the middle of the list. Thus, the isolated item in the middle of the list can be seen as being released from proactive interference that nonisolated items are exposed to. Fig. 3a shows that the isolated pattern at Serial position 9 has a larger net input than the nonisolated patterns. Notice that the degree of LTP is the same for all three conditions at the first serial position. Thus, LTP does not account for the fact that the isolation effect occurs at the first serial position.

3.2.5. Adaptive thresholds and LTD cause an isolation effect at the first serial position due to a lack of retroactive interference

LTD accounts for an isolation effect at the first serial position because a nonisolated item declines in performance. Cells representing the isolated item at the first serial position are likely to be inactive during encoding of the following nonisolated items. In these cells, there will be no LTD and no retroactive interference because the presynaptic cells are inactive. However, presynaptic cells representing a nonisolated item at the first serial position are more likely to be active during the presentation of the following nonisolated item yielding a retroactive interference. In these cells there will be LTD because the presynaptic cell is active. The stronger LTD in the nonisolated condition decreases the performance of Serial position 1 during encoding of the following serial positions so that the nonisolated items are less likely to be recalled compared to the isolated items (where there is less LTD).

The performance in Fig. 3a is higher for the isolated item at the first serial position and at Position 9. The adaptive threshold gives rise to a superior performance of isolated patterns over nonisolated patterns. At the first serial position the isolation effect occurs because of LTD, and at the following serial position it occurs due to a combination of LTP and LTD. The isolated pattern at Position 1 has a larger net input than the isolated pattern at Position 9 because the primacy effect contributes to the isolation effect at Position 1 but not at Position 9.

3.2.6. Adaptive thresholds yield less retroactive interference of an isolated pattern in comparison to a heterogeneous control list at all serial positions

Activity in an isolated pattern will cause less retroactive interference in synaptic strength than activity in a heterogeneous control list. In fact, active cells in the isolated pattern (that are not active in the correlated set of patterns) are only active when the isolated pattern is presented. Because these cells are very seldom active, the adaptive threshold will be close to zero, which makes the retroactive interference in the synaptic strength close to zero. In contrast, cells active in the heterogeneous list are active relatively more frequently, leading to more retroactive interference and worse performance than the isolated item.

The synaptic boundaries play an essential role in expressing this retroactive interference. The retroactive interference adds noise with a mean value of zero; however, without the bound-

aries the expected value of the net input is unaffected by this interference. The synaptic boundaries “squeeze” the expected value so that it decays toward zero at a rate that depends on the amount of retroactive interference.

3.3. Recall by relaxation

The performance in Fig. 3a was due to simplicity measured by the net input. Similar results are also obtained if performance is measured by the activity, where nodes are set to an active state if the net input exceeds a certain threshold (i.e., mathematically equivalent with a two-layer feedforward neural network). Here it is investigated how the model behaves when it is implemented by using relaxation in a recurrent network.

A simulation was conducted to confirm that the main results also hold when the activity pattern is recalled by relaxation (see Appendix A for details). To achieve reasonably successful recall, the number of nodes were increased to $N = 60$, the number of encoded patterns was decreased to 10, the weights were initialized to zero, and the minimum value of x was set to 1 (rather than zero).⁶ Simulations were run with adaptive thresholds and synaptic boundaries. A pattern was conceived as correctly recalled if the recalled activity exactly matched the encoded activity. The results in Fig. 3b are qualitatively similar to the case of no relaxation, and it accounts for all the data, which are primacy, recency, and isolation effects at all serial positions in comparison to heterogeneous and homogenous control lists. These results were obtained also when the threshold of activating nodes was varied; however, this threshold obviously influenced the overall performance.

3.4. Fits

Von Restorff (1933) conducted an experiment in which an isolated item was presented in a homogenous context, and a control condition consisted of a heterogeneous list of items. There were three experimental conditions corresponding to three lists. In List 1 there was one number and nine nonsense syllables, in List 2 there was one syllable and nine numbers, in List 3 all items were from different categories (i.e., one number, one syllable, one color, one letter, one meaningful word, one small photograph, one symbol, one button, one punctuation).

MIA was fitted to von Restorff’s data (1933) by using the explicit solution of the model as described in the Appendix. The mean squared error between fitted data and the experimental data was minimized by changing $N = 4.7$, $N_c = 3.2$, and $N_{act} = 1.4$ (i.e., the expected number of active nodes in each pattern). All other parameters were identical to the settings in the simulations.

The following results were obtained. The probability of recalling the isolated item was 0.70 in the experimental data and 0.70 in the simulated data. The probability of recalling the nonisolated items was 0.22 in the experimental data and 0.22 in the simulated data. The probability of recalling the isolated item in the heterogeneous control condition was 0.40 in the experimental data and 0.40 in the simulated data. The mean squared error was 0.00. Thus there was a complete agreement between empirical and fitted data.

3.5. Fixed adaptive threshold, unbounded synaptic strength, and a symmetrical learning rule do not account for the data

One goal of modeling behavior phenomena is to eliminate plausible alternative theories that cannot account for the behavioral data. Here it is briefly shown that fixed LTP–LTD thresholds, unbound synaptic strengths, and a symmetrical learning rule cannot account for the behavioral data. These results indicate that adaptive thresholds, synaptic boundaries, and an asymmetric learning rule are necessary to account for the empirical data in the MIA model. Unless specified otherwise these simulations are identical to the simulations in Fig. 3a. When no boundary was used, the learning rate was fixed to 1, and the weights were initialized to zero.

3.5.1. An adaptive threshold without boundaries eliminates three effects from the model (Fig. 4a)

First there is no isolation effect at the first serial position. Second, there is no isolation effect in comparison to a heterogeneous control list. Third, there is no power-function forgetting curve.

3.5.2. Fixed LTP–LTD threshold with boundaries eliminates the primacy effect and all the isolation effects

Fig. 4b. shows power-function forgetting curves when synaptic boundaries are used in combination with a fixed LTP–LTD threshold. The net input is larger for the correlated nonisolated

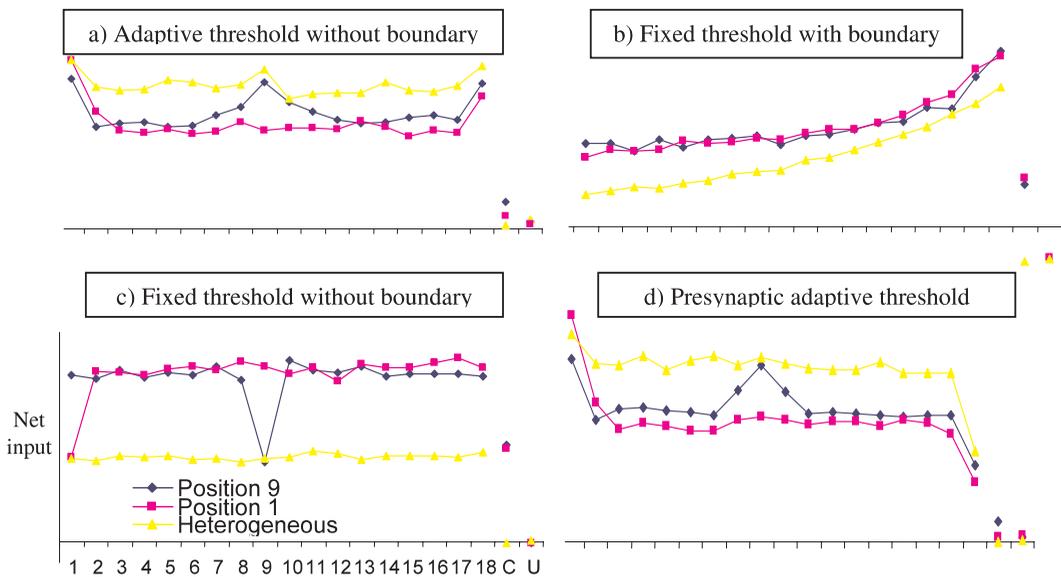


Fig. 4. Alternative models

Note. The panels are identical to Fig. 3a except that other rules for synaptic plasticity are used. Fig. 4a shows the prediction for an adaptive threshold without boundaries, Fig. 4b for a fixed threshold with boundaries, Fig. 4c for a fixed threshold without boundaries, and Fig. 4d shows the prediction for an adaptive threshold both at the presynaptic and postsynaptic cells without boundaries. See the text and the Appendix for details.

patterns compared to the heterogeneous control lists. This occurs because the threshold is inappropriately low for the correlated patterns ($a < p$), leading to an accumulation of synaptic strength during encoding of the correlated patterns.

3.5.3. Fixed LTP–LTD threshold without boundaries yields a negative isolation effect, no primacy, and no recency effect (Fig. 4c)

This simulation is identical to a standard Hopfield network (Hopfield, 1982,1984).

3.5.4. The recency effect occurs only if the adaptive threshold is introduced at the postsynaptic cell and not at the presynaptic cell

Many connectionist networks use symmetrical learning rules; for example, the Hopfield network is symmetric with a fixed threshold both at the pre- and the postsynaptic cell. In comparison, the learning rule in MIA is asymmetrical so that the strength of a connection in one direction may differ from the strength in the other direction. Fig. 4d shows what happens when a symmetrical learning rule is introduced in MIA so that the adaptive threshold is used both at the pre- and the postsynaptic cell. The main result is that no recency effect is found when an adaptive threshold is introduced in the presynaptic cell. Quite the contrary, the last items are recalled less well than items in the middle of the list—a negative recency effect. This negative recency occurs because the synaptic strength increases at the item following the to-be-recalled item. In particular the synaptic strength increases when adaptive threshold is high, and both the pre- and the postsynaptic cells are inactive. This demonstrates that the recency effect occurs, given that the adaptive threshold is introduced at the postsynaptic cell and is absent at the presynaptic cell.

4. Discussion

This article has shown that the isolation effect can be accounted for by applying empirical findings of synaptic plasticity in a neural network. An adaptive LTP–LTD threshold accounts for primacy, recency, and the isolation effect in the middle of the list. When an adaptive threshold is combined with synaptic boundaries, then an isolation effect is found at the first serial position, in comparison to the heterogeneous control list, and a power-function forgetting curve is obtained. These findings indicate the importance of basing memory models on empirical *measurable properties* of the neural system in the brain.

The adaptive threshold eliminates systematic drift in synaptic strength that otherwise occurs during encoding of correlated patterns. It causes LTP that enhances the synaptic strength during encoding of the first serial position, giving rise to a primacy effect and to an isolation effect in the middle of the list. LTD following encoding of the to-be-recalled item diminishes the performance of the nonisolated item or items, to a larger extent than the isolated item, causing an isolation effect at the first serial position. This may be compared with the novelty or salient accounts for the isolation effect that do not predict the isolation effect at the first serial position. LTP in combination with bounded synaptic strength also decreases the interference for the isolated item, which accounts for an isolation effect in comparison to a heterogeneous control list.

Without any additional assumptions, MIA also predicts the classical serial-position curve. LTP causes a primacy effect at the beginning of the list because the adaptive thresholds are low

at the beginning of the presentation of a list (i.e., the cells are assumed to be inactive before encoding). LTD following encoding of the to-be-recalled item also causes a recency effect. Furthermore, bounded synaptic strength with variability in learning rates predicts power-function forgetting curves. In summary, the results indicate that the isolation, the primacy, and the recency effects are accounted for by an adaptive threshold, whereas neither of these effects is found when the threshold is fixed.

The proposed mechanisms in the MIA model occur during encoding or storage, whereas no specific mechanism is needed at retrieval. This indicates that the accounted phenomena should occur for different tasks and across different retention intervals. These results support the view that short- and long-term memory share common principles rather than being separate memory systems (for an in-depth discussion of this topic, see, for example, Davelaar et al., 2005).

The isolation effect consists of a complex set of findings. It cannot simply be explained by novelty or salience because the phenomena also occurs if the isolated item is presented before the context has been established, for example, at the first serial position. Given an objective criteria of difference, it also seems to be difficult to account for the isolation effect by interference, because the isolated items show better performance than a control list of heterogeneous items in which each item is equally different from the isolated item. However, von Restorff (1933) originally argued for a relative measure of difference, where a continuous great difference in the heterogeneous control list creates an impression of a uniform whole, so an isolated item in a homogenous list is experienced as more distinct than an arbitrary item in a homogeneous control list. More recently Hunt and Lamb (2001) argued for a conceptualization of distinctiveness that acknowledges both similarities and differences among items. This conceptualization of distinctiveness was supported by data showing that the isolation effect is influenced by encouraging participants to process items via a difference judgment-orienting task.

Roberts and Pashler (2000, 2002) argued for three other criteria than goodness of fit for evaluating models, namely (a) the flexibility of theory (how much it cannot fit), (b) the variability of the data (how firmly the data rule out what the theory cannot fit), and (c) the likelihood of other outcomes (whether the theory can fit any plausible results). MIA also scores well on these criteria. (a) The theory constrains possible outcomes so that the primacy effect, the recency effect, and the isolation effect occur at any serial position in comparison to the heterogeneous list and a homogenous control list. These theoretical constraints are based on specific knowledge of synaptic plasticity and on computational constraints needed to encode correlated patterns. (b) The empirical data firmly agree with these constraints. That is, the primacy (see the review in the introduction [Section 1]), the recency, and the isolation effects are very robust phenomena and are consistently found under the conditions they are expected to be found. (c) Alternative possible outcomes of the empirical data would have been inconsistent with the theory. For example, a lack of an isolation effect (in comparison to either a heterogeneous or homogenous control list at any serial position), a lack of the primacy effect, or a lack of the recency effect would have been inconsistent with the theory. Beyond Roberts and Pashler's three criteria, this article has argued for additional, and at least as important, criteria for a good model, namely, a base model of established findings of neural plasticity.

By analyzing the explicit solution (provided in the Appendix), it is easy to realize that MIA makes the appropriate qualitative predictions for any allowed (or nondegenerated) choice of

parameters, given that the list length is sufficiently long to prevent interference from the recency and primacy effects (e.g., for a list length of two items either a primacy or a recency effect will occur depending on the choice of parameters). Furthermore, the number of free parameters in MIA is minimal. Like any distributed model, parameters are of course needed to describe the total number of nodes, the number of active nodes, and the degree of overlap between the nonisolated and the isolated patterns. In addition, there is only one parameter needed to specify the attenuation of the adaptive threshold (which is fixed to zero in this article). Parameters are needed to describe the learning rate distribution; however, the most essential parameter can be extracted from empirical data on the slope of the forgetting curves.

The adaptive threshold and LTP cause a larger strengthening of novel than familiar items. Several other successful learning rules also strengthen novel items more than familiar items. For example, the Rescorla–Wagner learning rule (Rescorla & Wagner, 1972) changes the weights, depending on the difference between the actual and the predicted associated strength of a conditional stimulus, and it has among other things been used to account for blocking in conditional learning. Other mathematical memory models have suggested novelty detectors (Lewandowsky, 1991; Lewandowsky et al., 1995; Metcalfe, 1993a; Simon & Lewandowsky, 2002); however, these models and the Rescorla–Wagner learning rule differ in several aspects from MIA. For example, the composite holographic associative recall model (CHARM; Metcalfe, 1993a, 1993b) was suggested because the composite trace in distributed systems can become catastrophically out of control. To solve this problem a monitoring system was introduced that measures the familiarity of the to-be-encoded item and adopts the learning rate so that novel items are encoded more than familiar items. It was argued that a breakdown in this familiarity system produces Korsakoff-like symptoms in the model, causing impairments in feeling-of-knowing judgments and a failure to release from proactive inhibition. CHARM's novelty detector uses a global measurement of familiarity and the mechanism implemented in this algorithm resembles the LTP mechanism in MIA. However, novelty detectors and the Rescorla–Wagner learning rule have no mechanism that resembles LTD, that is, no mechanism where a previously activated representation is more likely to decrease the synaptic strength. Therefore, novelty detectors, or the Rescorla–Wagner learning rule, cannot account for the isolation effect at the first serial position, nor can they account for the recency effect. Furthermore, CHARM's novelty detector is not directly related to synaptic plasticity.

In serial recall participants are required to recall items in the same order as they were presented. This version of MIA is not directly applied to serial recall. For example, the serial order of the encoded items is not explicitly represented in the network, and the obtained serial position curves are only mediated indirectly by the strength of the items. There is a large body of serial recall models; however, several of the key features in these models can be categorized into the following representational principles: (a) *Item marking* refers to association of the items with some independent representation of order such as time, temporal context, or list position (oscillator-based memory for serial order; Brown, Preece, & Hulme, 2000; Neath, 1999). For example, the start–end model (Henson, 1998) assumes that a start unit decreases and an end unit increases its activation from the start to the end of the list. Performance increases at the start and at the end of the list because the representation is more distinct. (b) *Primacy gradient* assumes that the strength of the encoding attenuates over serial position (Farrell

& Lewandowsky, 2002; Henson, 1996). (c) *Response suppression* makes the recalled item temporally unavailable for further report (Burgess & Hitch, 1999; Lewandowsky & Murdock, 1989); for example, participants have difficulty reporting both occurrences of repeated items (Duncan & Lewandowsky, 2005). (d) Recalling an item (*output*) *interferes* with the accessibility of the items yet to be recalled (Anderson & Neely, 1996). For example, Oberauer (2003) controlled for input and output order by appropriate cuing at test and found a primacy effect when the performance was plotted against output order as well as input order. (e) *Retroactive (input) interference* causes a recency effect because earlier items suffer interference from items presented later in the list (Nairne, 1988).

In contrast to the models referenced previously, MIA is based on neurophysiologic findings of synaptic plasticity. However, some, but not all, MIA mechanisms can be mapped into the classification schema listed previously. For example, MIA may be regarded as a primacy gradient model because the adaptive threshold diminishes LTP over list positions. The synaptic boundary causes retroactive input interference so that older items are recalled more poorly than new items. Item marking, response suppression, and output interference are not instantiated in this version of MIA. Furthermore, the blocking of LTD due to presynaptic inactivation in MIA is a mechanism that has not been introduced into earlier serial recall models. This LTD blocking is needed to account for the isolation effect at the first serial position and the isolation effect in comparison to a heterogeneous control list. It is therefore unclear whether current serial recall models can account for these aspects of the isolation effect.

High-priority events (HPE) is an experimental paradigm related to the isolation effect where participants are asked to pay special attention to, and to be sure to remember, specific target items in a list. Priority effect is often considered as yet another example of distinctiveness (e.g., Christianson & Loftus, 1987; Waugh, 1969). However, the negative effects of HPE are unlike the effects found with the standard isolation effect (which is based on physical or categorical distinctiveness). For example, HPE show a negative effect on the recall and recognition of preceding items for fast presentation times (i.e., 0.5–1 sec), although this effect is absent or less pronounced for long presentation times (i.e., 2 sec) and absent for or less pronounced for items following the HPE (Tulving, 1969). Tulving interpreted this as follows: For fast-presented items encoding occurs beyond the temporal interval that the stimulus is physically presented, and this encoding is interrupted by a new HPE, given that participants have been instructed to set a high priority for the HPE. MIA has not been applied to HPEs; however, it is straightforward to extend MIA to account for this phenomenon in the framework suggested by Tulving.

Basing cognitive and behavioral memory theories on established empirical knowledge of synaptic plasticity advances our understanding of cognition. This article has contributed to this formidable task by showing that the isolation effect and serial-position effect can be accounted for by using fundamental findings related to mechanisms in synaptic plasticity.

Notes

1. It is not immediately obvious which statistics to use to determine whether the PR is significantly different from zero, or to determine whether two PRs differ. It is problematic

to base the statistics on PRs calculated for each participant separately because the denominator may be zero or close to zero for some participants. One possibility is to assume binomial distribution with fixed probability correct for any serial position (or alternatively use the probabilities obtained in the experimental data). Monte Carlo simulations can then be conducted to obtain a 95% prediction interval and to evaluate whether the empirically obtained values lie within this interval (see, for example, Neter, Kutner, Nachtsheim, & Wasserman, 1996).

2. Notice that this way of measuring the last item in the primacy effect does not take into account whether the performance at each serial position is significantly different in a statistical sense. An alternative measurement of the last item in the primacy effect is the first item that is *not* statistically smaller than the preceding item. Although this measurement may intuitively seem more appealing, it suffers from the problem that it will bias the PR to a smaller value than the true value because the number of observations is limited.
3. More precisely, LTD occurs when the postsynaptic activity is less than the adaptive threshold but *larger* than zero. Thus, a moderately low degree of postsynaptic activity is required for LTD to occur. This model represents activity in two binary states (inactive or active), where the inactive state may be interpreted as a moderately low spike frequency (i.e., above zero but less than the adaptation threshold) where LTD can occur. A more biologically detailed model should of course use a continuous distribution of activity rather than a binary representation; however, we believe that this level of detail is sufficient to describe the phenomena of interest.
4. The important factor is that the expected *change* in synaptic strength is zero; however, the expected value of the synaptic strength may differ from zero (for example, depending on the activation threshold). This is believed to be a neurologically plausible assumption because it corresponds to the expectation that synaptic strength (over many synapses) does not change, or drift systematically, over long time periods.
5. Implementing the cue and the target patterns with identical patterns yields the same predictions as if these patterns were dissimilar (represented in two different populations of nodes) and randomly generated. This way of implementing retrieval is also similar to a recognition test. Furthermore, an implementation of free recall was avoided because it would require a substantially more complex retrieval schedule.
6. Experimental data on serial recall typically uses shorter list lengths than free recall data (e.g., six items in Oberauer, 2003). Simulations show that similar predictions are obtained for different list lengths.

Acknowledgments

I would like to thank John Wixted, Steve Gotts, Reed Hunt, Ben Murdock, Ausra Saudargiene, and Kenneth Malmberg for constructive comments on this article. Thanks to Andy Smart for correcting the English grammar. MIA is dedicated to love.

References

- Amit, D. J., & Tsodyks, M. (1991). Quantitative study of attractor neural networks retrieving at low spike rates: I. Substrate-spikes, rates, and neuronal gain. *Network: Computation in Neural Systems*, 2, 259–274.
- Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork & R. A. Bjork (Eds.), *Memory: Handbook of perception and cognition* (2nd ed., pp. 237–313). San Diego, CA: Academic.
- Atkinson, R., & Shiffrin, R. (1968). Human memory: A proposed system and its control processes. In K. Spence & J. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 2, pp. 89–195). New York: Academic.
- Bear, M. F. (1996). Progress in understanding NMDA-receptor-dependent synaptic plasticity in the visual cortex. *Journal of Physiology–Paris*, 90, 223–227.
- Bear, M. F., Cooper, L. N., & Ebner, F. F. (1987, July 3). A physiological basis for a theory of synaptic modification. *Science*, 237, 42–48.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1), 32–48.
- Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term memory free recall. *Cognitive Psychology*, 6, 173–189.
- Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, 107, 127–181.
- Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, 106, 551–581.
- Carlesimo, G. A., Marfia, G. A., Loasses, A., & Caltagirone, C. (1996). Recency effect in anterograde amnesia: Evidence for distinct memory stores underlying enhanced retrieval of terminal items in immediate and delayed recall paradigms. *Neuropsychologia*, 34, 177–184.
- Ceccarelli, B., & Hurlbut, W. P. (1980). Vesicle hypothesis of the release of quanta of acetylcholine. *Physiological Reviews*, 60, 396–441.
- Christianson, S.-Å., & Loftus, E. F. (1987). Memory for traumatic events. *Applied Cognitive Psychology*, 1, 225–239.
- Clothiaux, E. E., Bear, M. F., & Cooper, L. N. (1991). Synaptic plasticity in visual cortex: Comparison of theory with experiment. *Journal of Neurophysiology*, 66, 1785–1804.
- Cooper, L. N. (1995). *How we learn: How we remember: Toward an understanding of brain and neural systems: Selected papers of Leon N Cooper*. London, World Scientific.
- Cooper, L. N., Liberman, F., & Oja, E. (1979). A theory for the acquisition and loss of neuron specificity in visual cortex. *Biological Cybernetics*, 33(1), 9–28.
- Craik, F. I. M. (1970). The fate of primary memory items in free recall. *Journal of Verbal Learning and Verbal Behavior*, 9, 143–148.
- Craik, F. I. M., & Birtwistle, J. (1971). Proactive inhibition in free recall. *Journal of Experimental Psychology*, 91, 120–123.
- Crick, F. H. C. (1989, June 12). The recent excitement about neural networks. *Nature*, 337, 127–132.
- Dalezman, J. J. (1976). Effects of output order on immediate, delayed, and final recall performance. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 597–608.
- Darley, C. F., & Glass, A. L. (1975). Effects of rehearsal and serial list position on recall. *Journal of Experimental Psychology: Human Learning and Memory*, 104, 453–458.
- Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J., & Usher, M. (2005). The demise of short-term memory revisited: Empirical and computational investigations of recency effects. *Psychological Review*, 112, 3–42.
- Duncan, M., & Lewandowsky, S. (in press). The time course of response suppression: No evidence for a gradual release from inhibition. *Memory*.
- Einstein, G. O., McDaniel, M. A., & Lackey, S. (1989). Bizarre imagery, interference, and distinctiveness. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 137–146.

- Ellis, N. R., Detterman, D. K., Runcie, D., McCarver, R. B., & Craig, E. (1971). Amnesic effects in short-term memory. *Journal of Experimental Psychology*, *89*, 357–361.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin and Review*, *9*, 59–79.
- Fischler, I., Rundus, D., & Atkinson, R. C. (1970). Effects of overt rehearsal procedures on free recall. *Psychonomic Science*, *19*, 249–250.
- Frégnac, Y., & Shulz, D. (1994). Models of synaptic plasticity and cellular analogs of learning in the development and adult vertebrate visual cortex. *Advances in Neural and Behavioral Development*, *4*, 149–252.
- Froemke, R. C., & Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, *416*, 433–438.
- Gershberg, F. B., & Shimamura, A. P. (1994). Serial position in implicit and explicit tests of memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1370–1378.
- Glanzer, M. (1972). Storage mechanisms in recall. In G. H. Bower and J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 5, pp. 129–193). New York, Academic.
- Glanzer, M., & Cunitz, A. (1966). Two storage mechanism in free recall. *Journal of Verbal Learning and Verbal Behavior*, *5*, 351–360.
- Glenberg, A. M., Bradley, M. M., Kraus, T. A., & Renzaglia, G. J. (1983). Studies of the long-term recency effect: Support of the contextually guided retrieval hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 231–255.
- Gordon, P. C., Hendrick, R., & Levine, W. H. (2002). Memory load interference in syntactic processing. *Psychological Science*, *13*, 425–430.
- Green, T. (1956). Surprise as a factor in the von Restorff effect. *Journal of Experimental Psychology*, *52*, 340–344.
- Hebb, D. O. (2002). *The organization of behavior: A neuropsychological theory*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc. (Original work published 1949)
- Henson, R. N. A. (1996). *Short-term memory for serial order*. Unpublished doctoral dissertation, Cambridge University, Cambridge, England.
- Henson, R. N. A. (1998). Short-term memory for serial order: The start–end model. *Cognitive Psychology*, *36*, 73–137.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent computational abilities. *Proceeding of the National Academy of Sciences, USA*, *81*, 3088–3092.
- Hopfield, J. J. (1984). Neurons with graded responses have collective computational abilities. *Proceedings of the National Academy of Sciences USA*, *81*, 3008–3092.
- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*, 269–299.
- Hunt, R. R., & Elliot, J. M. (1980). The role of nonsemantic information in memory: Orthographic distinctiveness effects on retention. *Journal of Experimental Psychology: General*, *109*, 49–74.
- Hunt, R. R., & Lamb, C. A. (2001). What causes the isolation effect? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1359–1366.
- Hunt, R. R., & Mitchell, D. B. (1982). Independent effects of semantic and nonsemantic distinctiveness. *Journal of Experimental Psychology: Human Learning and Memory*, *8*, 81–87.
- Jahr, C. F. (2003). Drooling and stuttering, or do synapses whisper? *Trends in Neurosciences*, *26*, 7–9.
- Jenkins, W. O., & Postman, L. (1948). Isolation and spread of effect in serial learning. *American Journal of Psychology*, *61*, 214–221.
- Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory and Cognition*, *24*, 103–109.
- Katz, B. (1969). *The release of neural transmitter substances*. Springfield, IL: Thomas.
- Kelley, M. R., & Nairne, J. S. (2001). Von Restorff revisited: Isolation, generation, and memory for order. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 54–66.
- Laing, C. R., & Chow, C. C. (2002). A spiking neuron model for binocular rivalry. *Journal of Computational Neuroscience*, *12*, 39–53.
- Levy, W. B., Brassel, S. E., & Moore, S. D. (1983). Partial quantification of the associative synaptic learning rule of the dentate gyrus. *Neuroscience*, *8*, 799–808.

- Lewandowsky, S. (1991). Gradual unlearning and catastrophic interference: A comparison of distributed architectures. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 445–476). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Lewandowsky, S., & Li, S. (1995). Catastrophic interference in neural networks: Causes, solutions and data. In F. N. Dempster & C. J. Brainerd (Eds.), *Interference and inhibition in cognition* (pp. 329–361). San Diego, CA: Academic.
- Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, *96*, 25–57.
- Lewis, R. L. (1996). Interference in short-term memory: The magical number two (or three) in sentence processing. *Journal of Psycholinguistic Research*, *25*, 93–115.
- Lynch, M. (2003). Long-term potentiation and memory. *Physiological Reviews*, *84*, 87–136.
- Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997, January 10). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, *275*, 213–215.
- McArthur, L. Z. (1981). What grabs you? The role of attention in impression formation and causal attribution. In E. T. Higgins, C. P. Herman, & M. P. Zanna (Eds.), *Social cognition: The Ontario Symposium* (Vol. 1, pp. 201–246). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why are there complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning paradigm. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 24, pp. 109–164). San Diego, CA: Academic.
- McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1994). Illusory correlation in the perception of groups: An extension of the distinctiveness-based account. *Journal of Personality and Social Psychology*, *67*, 414–429.
- Metcalfe, J. (1993a). Monitoring and gain control in an episodic memory model: Relation to the P300 event-related potential. In A. F. Collins, S. E. Gathercole, M. A. Conway, & P. E. Morris (Eds.), *Theories of memory* (pp. 327–353). Hove, England: Lawrence Erlbaum Associates, Inc.
- Metcalfe, J. (1993b). Novelty monitoring, metacognition, and control in a composite holographic associative recall model. *Psychological Review*, *100*, 3–22.
- Miller, K. D. (1990). Correlation-based models of neural development. In M. A. Gluck, D. E. Rumelhart (Eds.), *Neuroscience and connectionist theory* (pp. 267–353). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Minsky, M. L., & Papert, S. A. (1988). *Perceptrons: An introduction to computational geometry* (Expanded ed.). Cambridge: MIT Press.
- Modigliani, V., & Hedges, D. G. (1987). Distributed rehearsals and the primacy effect in single-trial free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 426–436.
- Murdock, B. B. (1960). The distinctiveness of stimuli. *Psychological Review*, *67*, 16–31.
- Murdock, B. B. J. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*, 482–488.
- Nairne, J. S. (1988). A framework for interpreting recency effects in immediate serial recall. *Memory and Cognition*, *16*, 343–352.
- Nairne, J. S. (1990). A feature model of immediate memory. *Memory and Cognition*, *18*, 251–269.
- Neath, I. (1993). Contextual and distinctive processes and the serial position functions. *Journal of Memory and Language*, *32*, 820–840.
- Neath, I. (1999). Modelling the disruptive effects of irrelevant speech on order information. *International Journal of Psychology*, *34*, 410–418.
- Neter, J., Kutner, M. H., Nachtsheim, C. J., & Wasserman (1996). *Applied linear statistical models*. Chicago: Irwin.
- Nicoll, R. A., Kauer, J. A., & Malenka, J. A. (1992). The current excitement in long-term potentiation. *Neuron*, *1*(2), 97–103.
- Oberauer, K. (2003). Understanding serial position curves in short-term recognition and recall. *Journal of Memory and Language*, *49*, 469–483.
- O'Brian, E. J., & Myers, J. L. (1985). When comprehension difficulty improves memory for text. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 12–21.

- O'Reilly, R. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, 8, 895–938.
- Pezdek, K., Whetstone, T., Reynolds, K., Askari, N., & Dougherty, T. (1989). Memory for real-world scenes: The role of consistency with schema expectation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 587–595.
- Phillips, J. L., Shiffrin, R. M., & Atkinson, R. C. (1967). The effects of list length on short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 6, 303–311.
- Pillsbury, W. B., & Rausch, H. L. (1943). An extension of the Koler–Restorff inhibition phenomena. *American Journal of Psychology*, 56, 293–298.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: The effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical condition II: Current research and theory* (pp. 64–69). New York: Appleton-Century-Crofts.
- Rioutl-Pedotti, M. S., Friedman, D., & Donoghue, J. P. (2000, October 20). Learning-induced LTP in neocortex. *Science*, 290, 533–536.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing in psychology. *Psychological Review*, 107, 358–367.
- Roberts, S., & Pashler, H. (2002). Reply to Roders and Rowe (2002). *Psychological Review*, 109, 605.
- Rubin, D. C., Hinton, S., & Wenzel, A. (1999). The precise time course of retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 1161–1176.
- Rubin, D. C., & Wenzel, A. E. (1996). One hundred years of forgetting: A quantitative description of retention. *Psychological Review*, 103, 734–760.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986, October 9). Learning representation by backpropagating errors. *Nature*, 323, 533–536.
- Schmidt, S. R. (1991). Can we have distinctive theory of memory? *Memory and Cognition*, 19, 523–542.
- Sikström, S. (1999). Power function forgetting curves as an emergent property of biologically plausible neural networks model. *International Journal of Psychology*, 34, 460–464.
- Sikström, S. (2000). The TECO theory and lawful dependency in successive episodic tests. *Quarterly Journal of Experimental Psychology*, 53, 693–728.
- Sikström, S. (2001). The variance theory for the mirror effect. *Psychonomic Bulletin and Review*, 8, 408–438.
- Sikström, S. (2002). Forgetting curves: Implications for connectionist models. *Cognitive Psychology*, 45, 95–152.
- Sikström, S., & Jaber, M. Y. (2002). The power integration diffusion (PID) model for production breaks. *Journal of Experimental Psychology: Applied*, 8, 118–126.
- Simon, F., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin and Review*, 9, 59–79.
- Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3, 919–926.
- Steinbach, J. H., & Stevens, C. F. (1976). Neuromuscular transmission. In R. Linás & W. Precht (Eds.), *Frog Neurobiology*. Berlin, Germany: Springer-Verlag.
- Tan, L., & Ward, G. (2000). A recency-based account of the primacy effect in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1589–1625.
- Tsodyks, M. (2002). Spike-timing-dependent synaptic plasticity—The long road towards understanding neuronal mechanisms of learning and memory. *Trends in Neurosciences*, 25, 599–600.
- Tulving, E. (1969, April 4). Retrograde amnesia in free recall. *Science*, 164, 88–90.
- Wallace, W. P. (1965). Review of the historical, empirical, and theoretical status of the von Restorff phenomenon. *Psychological Bulletin*, 63, 410–424.
- Watkins, M. J., & Peynircioglu, Z. F. (1983). Three recency effects at the same time. *Journal of Verbal Learning and Verbal Behavior*, 22, 375–384.
- Waugh, N. C. (1969). Free recall of conspicuous items. *Journal of Verbal Learning and Verbal Behavior*, 8, 448–456.

- Wixted, J. T., & McDowell, J. J. (1989). Contributions to the functional analysis of single-trial free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 685–697.
- von der Malsburg, C. (1973). Self-organizing of orientation sensitive cells in the striate cortex. *Kybernetik*, 14, 85–100.
- von Restorff, H. (1933). Über die Wirkung von Bereichsbildungen im Spurenfeld. *Psychologische Forschung*, 18, 299–342.
- Wright, A. A., Santiago, H. C., Sands, S. F., Kendrick, D. F., & Cook, R. G. (1985, July 19). Memory processing of serial lists by pigeons, monkeys and people. *Science*, 229, 287–289.
- Zipser, D., & Andersen, R. A. (1988, February 25). A back propagation programmed network that stimulates response properties of a subset of posterior parietal neurons. *Nature*, 331, 679–684.

Appendix. A Neural Network Implementation of MIA

This Appendix shows how MIA can be implemented in a neural network related to a Hopfield network (Hopfield, 1982, 1984). By convention in the literature the synaptic strength is referred to as weights and cells as nodes in the Appendix. Node i at time t may either be in an active state represented as $\xi_i(t) = 1$ or an inactivate state represented as $\xi_i(t) = 0$. The weight change between node i and node j at time t , $\Delta w_{i,j}(t)$, is the product of activity of the postsynaptic activity minus the adaptive thresholds, $a_i(t)$, and the presynaptic activity. This weight change is scaled by the learning rate ($\eta_{i,j}$). The weight within a node is zero $w_{i,i}(t) = 0$:

$$\Delta w_{i,j}(t) = \eta_{i,j}(\xi_i(t) - a_i(t))\xi_j(t) \text{ for } i \neq j \text{ and otherwise } w_{i,i}(t) = 0$$

In Fig. 4d an alternative (nonstandard) learning rule is simulated in MIA, where an adaptive threshold is introduced in both the presynaptic and the postsynaptic node:

$$\Delta w_{i,j}(t) = \eta_{i,j}(\xi_i(t) - a_i(t))(\xi_j(t) - a_j(t)) \text{ for } i \neq j \text{ and otherwise } w_{i,i}(t) = 0$$

The implementation is more similar to the Hopfield model (i.e., it is identical given that the threshold is fixed) and it also makes the synaptic strength symmetrical (i.e., $w_{i,j} = w_{j,i}$), which is important for reaching stable states during relaxation in Hopfield networks.

The adaptive threshold is set to a weighted average of earlier activity, where more recent activities are weighted heavier than older activities. This is implemented by setting the (new) adaptive threshold at time t to the activity at the previous time step $t - 1$ (weighted with $1 - c$) plus the previous adaptive threshold (weighted by c):

$$a_i(t) = c a_i(t - 1) + (1 - c)\xi_i(t - 1), 0 \leq c < 1$$

where the constant c determines the size of the time window that the activity is averaged. For example, $c = 0$ sets the adaptive threshold to the activity at the previous time step, which is used in this article. For larger c the activity is weighted over several time steps back, where the more recent time steps are more heavily weighted.

The adaptive thresholds are initialized to zero prior to encoding of the first item. The model may be simulated with weight boundaries and a constant learning rate; however, this would produce an exponential rate of forgetting. A forgetting curve that approximates a power-fun-

ction requires a distribution of learning rates. At the start of the simulation of each participant, the learning rates are generated from a power-function distribution

$$\eta_{i,j} = \frac{1}{x^{\alpha-1}},$$

where x is a random number with a uniform distribution ($0 < x < x_{max}$) and $\alpha > 0$ is a constant that determines the slope of the forgetting curve (Sikström, 1999; 2002). The larger x_{max} is the larger retention interval the model will show a good approximation to a power-function forgetting curve. However, if x_{max} is too large then most of the weight changes will be very small, leading to a poor overall performance in the network (Sikström, 1999; 2002). Therefore it is useful to set x_{max} to a moderate large value.

Each weight is initialized to a random value between $-b$ and b . The weight at time t equals the weight at time $t - 1$ plus the weight change at time t . Furthermore, the weights are bounded so they cannot exceed an upper boundary (b) and not decrease below a lower boundary ($-b$) ($w_{i,j}(t) = \text{Max}[\text{Min}[w_{i,j}(t-1) + \Delta w_{i,j}(t), b], -b]$)

The performance is measured by the net inputs to active nodes. The net input to node i equals the sum of all weights connecting node i to active nodes:

$$h_i(t) = \sum_{j=1}^N w_{i,j}(t) \xi_j(t)$$

An explicit solution of the MIA model

This Appendix calculates an explicit analytical solution of MIA.

Let N_{act} represent the expected number of active nodes in a pattern, $p = N_{act}/N$ the probability that a node is active among the N nodes in the heterogeneous list, and $p = N_{act}/N_c$ the probability that a node is active among the N_c number of nodes representing the nonisolated items. Let a_k represent the adaptive threshold for an item in serial position k . The sliding modification is initialized to zero ($a_1 = 0$) and is set to $a_{k+1} = a_k c + (1 - c) p$, for $k > 1$.

Let w_k represent the expected weight at serial position k . LTP causes the following increase in weight for nonisolated (and heterogeneous) items $w_k = (1 - a_k)$, and for isolated items $w_k = (1 - N_c/N a_k)$.

LTD occurs following encoding of an item and it decreases the expected weight change. Let $p' = N_{act}/N$ for an isolated item and $p' = p$ for a nonisolated item. Let $a'_{k+1} = a_k c + (1 - c)$ be the adaptive threshold for item k at serial position $k + 1$, and $a'_{k+1} = a'_k c + (1 - c) p'$ be the adaptive threshold for item k at serial position $k + 2$ or larger. The expected weight for item k when tested at serial position K following LTD is

$$w_k = w_k + \sum_{i=k+1}^K (p' - a'_i) p'$$

The weight boundaries cause attenuation in the weights over the retention interval so that the weight has attenuated to w'_K at serial position $K \geq k$. The expected absolute weight change

is $2p(1-p)$, the learning rate $x^{-1/(1-\alpha)}$, and the distance between the boundaries $2b$. Single weights attenuate exponentially with an exponent of $K-k$, and the expected attenuation can be expressed as:

$$w'_K = w_k \frac{1}{x_{\max}} \int_0^{x_{\max}} \left(1 - \text{Min} \left(1, 2p(1-p)x^{-\frac{1}{1-\alpha}} \frac{1}{2b} \right) \right)^{K-k} dx$$

For nonisolated items the expected summed net input for all nodes encoded in the active is

$$h_k = N_{act}^2 w'_k$$

For isolated items, only a subset of the nodes attenuates:

$$h_k = N_{act}^2 \left(\frac{N_c^2}{N^2} w'_k + \left(1 - \frac{N_c^2}{N^2} \right) w_k \right)$$