Super Memory Bros.: Going from mirror patterns to concordant patterns via similarity enhancements

JASON D. OZUBKO

University of Waterloo, Waterloo, Ontario, Canada

AND

STEVE JOORDENS

University of Toronto, Scarborough, Toronto, Ontario, Canada

When memory is contrasted for stimuli belonging to distinct stimulus classes, one of two patterns is observed: a *mirror pattern*, in which one stimulus gives rise to higher hits but lower false alarms (e.g., the frequency-based mirror effect) or a *concordant pattern*, in which one stimulus class gives rise both to higher hits and to higher false alarms (e.g., the pseudoword effect). On the basis of the dual-process account proposed by Joordens and Hockley (2000), we predict that mirror patterns occur when one stimulus class is more familiar and less distinctive than another, whereas concordant patterns occur when one stimulus class is more familiar than another. We tested these assumptions within a video game paradigm using novel stimuli that allow manipulations in terms of distinctiveness and familiarity (via similarity). When more distinctive, less familiar items are contrasted with less distinctive, more familiar items, a mirror pattern is observed. Systematically enhancing the familiarity of stimuli transforms the mirror pattern to a concordant pattern as predicted. Although our stimuli differ considerably from those used in examinations of the frequency-based mirror effect and the pseudoword effect, the implications of our findings with respect to those phenomena are also discussed.

In typical recognition memory experiments, participants are asked to discriminate items they have seen previously during a study phase (i.e., “old” items) from those they have not (i.e., “new” items). Classifying an old item correctly is called a *hit*, and mistakenly classifying a new item as old is called a *false alarm*. The difference between hits and false alarms provides an indicator of memory strength, and to assess effects on memory, differences can be contrasted across stimulus categories or empirical manipulations.

The present article focuses primarily on differences between stimulus classes; within the memory theory literature, two *patterns of effects* have informed memory theories. In cases where Stimulus Class A gives rise to both greater hits and lower false alarms than does Stimulus Class B the resulting pattern generally is called a *mirror pattern*, because the advantage of Class A in terms of more hits is mirrored in terms of fewer false alarms to that class. However, when Stimulus Class A gives rise to both higher hits and higher false alarms relative to Class B, the result is called a *concordant pattern*, because the hits and false alarms move in the same direction across the stimulus classes. The primary goal of the present work is to gain a better understanding of the factors that underlie the emergence of mirror patterns in some contexts and concordant patterns in others.

To aid in this goal and to provide a more explicit theoretical starting point for our considerations, we focus on a specific mirror effect and on a specific concordant effect, and we consider the potential factors relevant to each. With respect to the mirror pattern, the *frequency-based mirror effect* is the finding that low-frequency words give rise to more hits and fewer false alarms than do high-frequency words (e.g., Glanzer & Adams, 1985). With respect to the concordant pattern, the *pseudoword effect* is the finding that pseudowords (i.e., nonwords or very rare words) elicit more hits and more false alarms than do words (Greene, 2004; Hintzman & Curran, 1997; Whittlesea & Williams, 2000; Wixted, 1992). Although typically they are considered separately, we consider these two phenomena conjointly in order to assess the theoretical parallels and differences.

Both single- and dual-process accounts have been proposed to account for the frequency-based mirror effect. Traditional single-process accounts (e.g., Gillund & Shiffrin, 1984; Hintzman, 1988; Murdock, 1982) suggest that for each type of stimulus, an *old* and a *new* distribution can be placed on a scale of memory strength, on which participants adopt a criterion above which they call an item to be old. The frequency-based mirror effect often is viewed as a challenge to these models, because, to account

for the mirror pattern, newly introduced low-frequency words must be assumed to have less average strength relative to high-frequency words, but to have more strength on average when they are old. Essentially, one needs to explain why presenting an item at study would give disproportionate advantage to low-frequency items to such an extent that their strength distribution would not merely get closer to, but would leapfrog, the strength distribution for high-frequency items (Glanzer & Adams, 1985).

More recent variants of single-process models have been able to account for the mirror effect by assuming that the *old* decisions are based on either a likelihood ratio or on a Bayesian probability calculation assessing the likelihood of an item being old versus new (e.g., Hilford, Glanzer, & Kim, 1997; Shiffrin & Steyvers, 1997). However, although these models can reproduce the basic mirror pattern, they have not yet provided a principled theoretical account of the fluctuations in the basic mirror pattern that occur when one manipulates those variables that typically are associated with familiarity and recollective memory influences (e.g., Balota, Burgess, Cortese, & Adams, 2002; Hirshman, Fisher, Henthorn, Arndt, & Passannante, 2002; Joordens & Hockley, 2000; Reder et al., 2000).

Dual-process accounts of the mirror effect (e.g., Joordens & Hockley, 2000; Reder et al., 2000; see Yonelinas, 2002, for a review) assume that recognition performance in general, and the frequency-based mirror effect in particular, reflect the combined effects of two different influences of memory: familiarity and recollection. Specifically, in the dual-process account proposed by Joordens and Hockley, high-frequency words are assumed to give rise to more familiarity than are low-frequency words, given that participants have more preexperimental experience with those items. Taking into account familiarity only, one would expect a greater degree of hits and false alarms to high- versus low-frequency words. However, it is assumed further that low-frequency words are bound more distinctively to the study list context, given that they have fewer preexperimental contextual associations and that this higher level of distinctiveness makes them easier to recollect consciously when they are old (e.g., Geraci & Rajaram, 2002; Schmidt, 1991; Valentine, 1991). Thus, the false alarm portion of the mirror effect is viewed as simply reflecting the higher familiarity of high- relative to low-frequency items, supporting more false alarms to high-frequency new items relative to low-frequency new items. The hit rate portion is viewed as reflecting situations in which the enhanced recollection of low-frequency items is sufficiently strong to overcome the higher familiarity of high-frequency words, supporting an overall tendency to call low-frequency words old more often when they are old.

Results from Joordens and Hockley (2000) and Reder et al. (2000) are often touted as providing strong support for the dual-process view described above. This is because they demonstrate that a range of manipulations that primarily should target one's ability to recollect (e.g., divisions of attention, manipulations of study-test lag, reductions in study time or response time at test) all lead to dramatic and predictable changes in the mirror pattern.

In contexts where the ability to recollect is decreased, the old portion of the mirror effect can be eliminated or even reversed (i.e., higher hits to high-frequency than to low-frequency items) while either having no effect on the false alarm portion or even enhancing the typical pattern. These results follow completely from the dual-process view and have been confirmed independently via different manipulations of ability to recollect, including manipulations such as administration of Midazolam to induce temporary loss of recollective memory (Hirshman et al., 2002) and comparisons of patient populations who have more or less severe damage to their recollection processes (Balota et al., 2002).

Recently, however, Malmberg, Zeelenberg, and Shiffrin (2004) demonstrated that the single-process, REM (retrieving effectively from memory) model of memory (Shiffrin & Steyvers, 1997) can provide reasonable fits to Hirshman et al.'s (2002) data. This result highlights the fact that, although dual-process models make straightforward predictions about how recognition memory should change under different conditions, single-process models can sometimes be fit to the data in a post hoc manner. Given such a result, it cannot be said that single-process models truly provide no account for described fluctuations of the mirror effect; however, it is reasonable to say that they do not provide a principled explanation for why such changes occur. We will return to single-process interpretations of our findings in the General Discussion.

Thus, taking the dual-process account of the mirror effect as our starting point, we can assume that at least two factors are critical for observing a mirror pattern. One stimulus class must be more familiar than another, but the less familiar stimulus class must also be more distinctive, thereby allowing greater conscious recollection. With these factors in mind, we turn to a consideration of concordant patterns via an examination of the current theories concerning the pseudoword effect.

The pseudoword effect is the finding that hit and false alarm rates are typically higher for pseudowords compared with words (Greene, 2004; Hintzman & Curran, 1997; Hockley & Niewiadomski, 2001; Whittlesea & Williams, 2000; Wixted, 1992). Like the mirror effect, the pseudoword effect also can be explained in terms of familiarity, if one assumes that pseudowords give rise to more familiarity than do words (Greene, 2004). This enhanced familiarity leads to both higher hit rates and higher false alarm rates for pseudowords compared with words. Support for the notion that pseudowords are more familiar than words comes from examining, in detail, the "old" responses to words and pseudowords. "Old" responses can be grouped into "remember" and "know" responses, which are assumed to provide rough indexes of the influences of familiarity and recollection, respectively (Gardiner & Java, 1993; Joordens & Hockley, 2000; although see Rotello & Macmillan, 2006; Wixted, 2007; and subsequent discussions in this article for some debate on the topic). If pseudowords are more familiar than words (which is indeed the case), one would therefore expect more "know" responses for pseudowords than for words (Gardiner & Java, 1990; Greene, 2004).

However, the assertion that pseudowords give rise to more familiarity than do words should seem odd, especially considering that participants seldom if ever have experienced pseudowords prior to the experiment, whereas words have been experienced repeatedly. Critically, though, preexperimental experience is not the only source for familiarity, which can arise due to a number of factors that are assumed to affect fluency of processing. In addition to preexperimental experience (e.g., Jacoby & Dallas, 1981; Joordens & Hockley, 2000; Reder et al., 2000), such factors include recent exposure to the stimulus (e.g., Higham & Vokey, 2000; Jacoby & Whitehouse, 1989; Joordens & Merikle, 1992) and similarity to other items in the experiment (e.g., Gillund & Shiffrin, 1984; Greene, 2004; Hintzman, 1988), which, according to Greene, is the most likely culprit for the pseudoword effect.

Previous studies have found that participants are more likely to respond with a false alarm to items that are similar to those seen recently (e.g., at study) than to less similar items. For example, Wickelgren (1966) found that if the number 12 had been seen at study, then a “reversed” foil (e.g., 21) would be disproportionately likely to elicit a false alarm, as would a foil that retained an element from the original stimulus (e.g., 17 or 24). More recently, Roediger and McDermott (1995) found that, after presenting a series of words all related to an underlying critical word (e.g., SLEEP), participants were very likely to believe they had also seen the critical word, even though it was never presented at study—in fact, frequently participants even falsely reported recollecting highly familiar words, leading to false alarms in which participants indicated they “remembered” that the item was old, and did not just “know” it was (see Deese, 1959). Recently, Cleary, Morris, and Langley (2007) have argued that, whereas structurally regular novel items are sometimes believed to give rise to increased false-positive responses (e.g., Whittlesea & Williams, 1998), this is true only when structural regularity is confounded with similarity to studied items. Hence, similarity to study, both at an orthographic and at a conceptual level, seems able to increase subjective feelings of familiarity and false-positive responses, and sometimes dramatically so.

In order to explain the higher familiarity of pseudowords relative to words in terms of similarity, one must assume that pseudowords are more similar to one another than words are to one another. Greene (2004) provides a reasonable rationale for such an assumption: Whereas words can be discriminated in terms of both orthography and semantics (as well as other dimensions, but we will focus on these two), pseudowords likely give rise to little, if any, semantic activation. Because semantics can be used to discriminate even highly orthographically similar words (e.g., HOUSE vs. HORSE), by lacking semantics, pseudowords are potentially more difficult to discriminate from one another (e.g., GLAWK vs. GRAWK). Hence, the lack of an extra discriminable dimension might make pseudowords more similar to one another than are words. This similarity gives rise to more familiarity for pseudowords than for words, which in turn results in increased hits and false alarms—that is, the *pseudoword effect* (see

Greene, 2004; Joordens, Ozubko, & Niewiadomski, 2008, for data supporting this contention; see Wixted, 1992, for a similar explanation regarding concordant effects with rare words).

Returning to the primary goal of the present work, the mirror effect and the pseudoword effect provide examples of recognition memory phenomena that are unified readily under the dual-process account provided above. According to this framework, a mirror pattern should arise when one stimulus class is both more familiar and less distinctive than the other, and a concordant pattern should arise as familiarity increases for one stimulus class but not the other. From our analysis of the pseudoword effect, we can also assume that, by increasing interitem similarity, one could increase the familiarity of a stimulus set. The experiments that follow attempt to assess these general principles derived from the memory literature in a novel context to determine whether they can provide a general description of the factors that would give rise to a mirror versus a concordant pattern.

We chose not to use words or pseudowords as our stimuli for two reasons. First, given all the dimensions on which such stimuli can vary, manipulating similarity and/or distinctiveness of words and pseudoword can be difficult. Second, we wished to test the notion that such phenomena as the mirror and pseudowords effects are generalizable memory phenomena and not simply artifacts of linguistic research or methodology. That is, researchers investigating phenomena such as the mirror and pseudoword effects make claims about the ways in which such stimulus features as similarity and distinctiveness affect familiarity and recollection. As a direct test of these hypotheses, we created a set of stimuli that varied systematically along these dimensions. Thus, if these theories are correct, despite not using words or pseudowords, mirror and concordant patterns should still arise predictably.

As an additional test of the generalizability of memory phenomena, we used both a traditional serial list study phase and a more ecologically valid study phase (i.e., a video game study phase). That is, we used cartoon-like characters as our stimuli and had participants study these characters, either in a typical serial list study phase or embedded within a video game context. Although video games themselves do not mimic the conditions of the real world, we feel they are more realistic in a cognitive sense. In video games, as in real life, the relevant information is embedded within a dynamic environment, and to perform successfully in the overall task, participants must pay attention to more than just the target items amid an abundance of useless and irrelevant information.

EXPERIMENTS 1A AND 1B

Replicating the Mirror Effect in Two Paradigms

Given our decision to use cartoon-like video game characters instead of words, the first step was to create two sets of characters that would give rise to a mirror pattern in both a traditional old/new recognition paradigm (Experiment 1A) and a video game recognition paradigm (Experiment 1B). To this end, we created two sets of characters varying with

respect to familiarity and distinctiveness that were used in both experiments (see Figure 1). The first set, normal characters, was simply a baseline set of cartoon-like characters from which to create the second set, which was a highly blurred and colored version of the first. The goal was to both enhance interitem similarity (thereby increasing familiarity) and reduce distinctiveness (thereby decreasing recollection) relative to the normal characters. That is, although the most distinctive features and color patterns were lost, we felt that the shapes were different enough for the characters to retain a unique identity. However, blurring gave rise to the side effect of increased similarity. That is, all the blurred characters looked similar to one another inasmuch as they all resembled some kind of odd blob.¹

With respect to producing our mirror pattern, the blurred characters could be seen as analogous to high-frequency words, because they should give rise to less recollection and to more familiarity than would normal characters. Normal characters then could be seen as analogous to low-frequency words, because they should be more distinctive and less similar to one another than would blurred characters. Note that the familiarity attributed to high-frequency words likely results from preexperimental experience and not from above-average similarity to one another, and so, our blurred characters may not be entirely analogous to high-frequency words. However, because both preexperimental experience and similarity can give rise to familiarity, and because familiarity is the factor we wish to increase, we felt that the blurred characters would be

satisfactory by at least approximating the effects observed with high-frequency words. Thus, it was not so much our goal to recreate the frequency-based mirror effect as it was to recreate the same interplay of underlying influences as is thought to underlie the frequency-based mirror effect by using the dual-process account we described earlier.

Method

Participants. Participants were recruited from the University of Toronto, Scarborough, and participated for one half hour to earn a 0.5 bonus credit toward their introductory psychology course. Sixty-two participants were in Experiment 1A, and 24 were in Experiment 1B. In Experiment 1A, eight outliers had hits or false alarm values that were at least 2 *SDs* above or below the mean values. Including these outliers in the analyses resulted in the false alarm difference approaching, but not reaching, significance [$t(62) = 1.88, p = .065$]; without the outliers, this difference was highly significant (as will be discussed in more detail shortly). Because removal of these outliers did not change the direction of any trends or qualitatively affect the outcome of any other results, but did result in the false alarm difference becoming highly significant, these outliers were dropped from the analyses. For completeness, however, we report both the adjusted and unadjusted means and *SDs* for Experiment 1A. The full data is presented in italics in Table 1, although we will focus on the unitalized data, from which outliers are excluded.

In Experiment 1B, three outliers had hits or false alarm values that were at least 2 *SDs* above or below the mean values. However, removing these outliers did not result in any qualitative change in the significance patterns observed; therefore, we left them in the analyses for completeness. All participants had normal or corrected-to-normal vision, and all were sufficiently proficient in English to complete the task without difficulty.

Procedure. Participants were told they would be taking part in a brief memory experiment. After signing a consent form, each participant was seated in a separate room with a computer. In Experiment 1A, at study, 40 characters (20 normal and 20 blurred) were intermixed randomly and appeared on the screen one at a time for 1 sec with a 0.5-sec interstimulus interval. In Experiment 1B, participants played through a short "Super Mario"-like game, during which they encountered the same characters as in Experiment 1A, but within the video game context.

The video game consisted of a simple, 2-D, side-scrolling game. Participants navigated a character that could walk left or right and jump through a game world consisting of pits, blocks, and other obstacles. As the participants traveled further to the right, the game world scrolled on. Throughout the world, they encountered the normal and blurred characters. The characters would walk or fly about, sometimes aimlessly, and the participants had to get close to the characters and then jump on them to defeat them. Every character had to be defeated before the study phase would end, and participants were informed of this fact. The world was also populated with random deposits of coins and power-ups that the player could collect for bonus points (the points were irrelevant in terms of the experiment).

Additionally, participants who contacted a character, but did not jump on it, would become injured or defeated. If they were defeated, the game paused for a few seconds, played a defeat sound, then a new player-controlled character would appear and the participant could continue to play. If only injured, the participant's character would shrink. Participants could cure their characters by getting one of the power-ups found in the game world. After playing through the game world, participants had to fight a boss in the game. This served as a temporary distraction to eliminate any recency effects. The boss fight lasted 5 min or until the boss was defeated.

In both Experiments 1A and 1B, participants were told to try to remember the characters as best they could, because they would later be tested to see how well they remembered them. At test, in both experiments, the study characters were shown again randomly

Figure 1. Examples of character types used in the experiments. Each blurred character is juxtaposed with its normal source character, making the similarities between each pair more readily apparent. Note however that, due to the sheer number of characters encountered in Experiments 1A and 1B (40 at study and 80 at test) and to the fact that normal and blurred versions of the same character rarely would be presented sequentially, the chance of participants noting analog characters is reduced greatly, if not eliminated.

Table 1
Means (and Standard Deviations) for Hits and False Alarms
for Different Character Types in Experiments 1A–3B

Experiment	Study Phase	Hits				False Alarms			
		<i>M</i>		<i>SD</i>		<i>M</i>		<i>SD</i>	
		Normal		Blurred		Normal		Blurred	
1A ^a	Serial list	.71	.15	.64	.15	.30	.14	.37	.18
		.71	.17	.64	.17	.33	.16	.38	.19
1B	Video game	.63	.15	.51	.19	.26	.13	.38	.20
2A	Video game	Normal		Multifeature		Normal		Multifeature	
		.62	.13	.55	.19	.26	.12	.35	.15
3A	Serial list	.72	.12	.68	.13	.35	.16	.44	.16
2B	Video game	Normal		Minimal Feature		Normal		Minimal Feature	
		.58	.16	.68	.19	.22	.13	.63	.23
3B	Serial list	.64	.14	.80	.16	.23	.13	.65	.16

^aNonitalicized values are those calculated after excluding eight outliers whose hits or false alarms deviated by at least 2 *SDs* from the means. Italicized values are those calculated from the complete data set for Experiment 1A.

intermixed with an equal number of new items. Participants were told to indicate which items were old by pressing the “/” key and which were new by pressing the “Z” key and to do so as quickly, but as accurately, as possible. After the experiment, participants were debriefed as to the nature of the study.

Materials and Apparatus. A set of 40 cartoon-like “normal” characters, similar to those found in the Super Mario series of Nintendo games (www.nintendo.com), were created for this experiment. Examples are shown in Figure 1. These characters tended to resemble mythical creatures or different animal species and were assumed to be a fairly distinctive set of items because each had many unique features and aspects. Normal characters were used as a baseline from which additional sets of characters were created.

A second set of characters was created by heavily pixelating (digitally blurring) and colorizing the 40 normal characters. Blurring was accomplished by graphically averaging together the nearest few pixels of every fourth pixel, which resulted in a substantial loss of detail in the characters. Furthermore, all blurred characters also were colorized with a random color and then dulled. Thus, the final characters were blurry and dull colored compared with the normal characters from which they were derived. This new blurred set was assumed to be far less distinctive than the normal characters because most distinguishing features and colors had been eliminated effectively, making these altered versions more similar to one another because each character was somewhat amorphous. However, because each blurred character maintained a relatively unique outline compared with other characters, we assumed that each would have at least one distinctive characteristic.

Results and Discussion

The results of Experiments 1A and 1B are shown in Table 1. Again, for Experiment 1A we analyzed the unitalicized data, which excluded the eight outliers. Both experiments were analyzed separately in a 2 (old vs. new) $\times$ 2 (normal vs. blurred characters) within-subjects ANOVA. In Experiment 1A, no main effect was found for character type [$F(1,54) < 1$]; however, participants could discriminate old from new [$F(1,54) = 359.00$, $MS_e = 0.02$, $p < .01$, $\eta^2 = .87$], and this factor interacted with character type [$F(1,54) = 27.48$, $MS_e = 0.01$, $p < .01$, $\eta^2 = .34$]. Paired-sample *t* tests revealed the presence of a full mirror pattern: Hit rates were higher for normal than for blurred characters [$t(54) = 2.56$, $p < .05$], and false

alarm rates were higher for blurred than for normal characters [$t(54) = 2.71$, $p < .01$].

In Experiment 1B, no overall difference between normal and blurred characters was found [$F(1,23) = 0$, $p = 1.00$]. However, participants could discriminate old from new items [$F(1,23) = 72.40$, $MS_e = 0.02$, $p < .01$, $\eta^2 = .76$], and this factor again interacted with character type [$F(1,23) = 31.20$, $MS_e = 0.01$, $p < .01$, $\eta^2 = .58$]. Paired-sample *t* tests confirmed that a full mirror pattern was obtained: Participants recorded more hits [$t(23) = 3.29$, $p < .01$] and fewer false alarms [$t(23) = 3.23$, $p < .01$] to normal characters than to blurred characters.

In both Experiment 1A and 1B, we replicated a full mirror pattern by using novel characters. As predicted, the more distinctive normal characters elicited a higher hit rate than did the blurred characters, whereas the generally more similar blurred characters elicited a higher false alarm rate than did the normal characters. Moreover, the mirror pattern seemed somewhat exaggerated in the video game context compared with the laboratory context, because larger effect sizes were noted.

Some single-process accounts of the mirror effect embody the assumption that a mirror effect arises whenever one set of stimuli is less memorable or discriminable than another (Glanzer & Adams, 1985, 1990). In our case, the blurred characters were less discriminable than the normal characters, and so these results are consistent with this type of single-process account as well as our dual-process account.

EXPERIMENTS 2A AND 2B

Playing With Similarity

Although a mirror pattern was replicated in Experiments 1A and 1B, it remains possible that familiarity and distinctiveness have nothing to do with mirror patterns, and that we just happened to select stimuli that gave rise to a mirror pattern for some other reason. In order to confirm that distinctiveness and familiarity are the driving factors behind

the mirror pattern in these experiments, in Experiment 2A we used a radically different method for generating a second set of characters to compare with the normal characters: Instead of blurring normal characters, we generated novel characters with a variety of distinctive features. We termed them *multifeature characters* (see Figure 1).

Multifeature characters were created out of different-shaped and -colored blocks. Without distinctive characteristics, this class of characters would no doubt be highly similar to one another. To make these characters distinctive, we added unique features such as mouths, eyes, and limbs, despite which, we still predicted that these characters would be less distinctive than normal characters. That is, although normal characters have numerous unique or unusually distinctive characteristics that set them apart from other members of their stimulus class, multifeature characters have only one or two distinctive features. However, because multifeature characters were still all noticeable as simple shapes and composed of a limited number of colors that were reused for many other multifeature characters, we predicted that multifeature characters would be more similar to other members of their stimulus set than normal characters would be to theirs. This interitem similarity should give rise to more familiarity for multifeature characters than for normal characters. Therefore, even though these multifeature characters appear quite different from the blurred characters in Experiments 1A and 1B, we predicted that they would give rise to a mirror pattern because they varied in terms of familiarity and distinctiveness as the dual-process account predicts they should to produce a mirror pattern.

Recall the dual-process prediction that mirror patterns occur when one set of stimuli is both more familiar and less distinctive than another, whereas concordant patterns occur when one set of stimuli are simply more familiar than another. In Experiment 2B, we test, this assertion by systematically reducing the distinctiveness of the multifeature items used in Experiment 2A, essentially reducing the number of dimensions on which they can be discriminated and leaving only the features that are similar between them. Thus, as we removed distinctive features and retained the common ones, we predicted that interitem similarity necessarily would increase. Studies focusing on words have found that increasing interitem similarity can lead to increases in both hits and false alarms (Criss & Shiffrin, 2004; Shiffrin, Huber, & Marinelli, 1995; Sommers & Lewis, 1999; Zaki & Nosofsky, 2001), and thus we predicted that a concordant pattern should emerge in Experiment 2B.

Finally, although the mirror pattern was observed previously in both a traditional (Experiment 1A) and a video game (Experiment 1B) study session, it was found more reliably with a video game paradigm. That is, several outliers were noted in Experiment 1A, which resulted in borderline false alarm differences. Because it seemed to be the more reliable paradigm to use with our stimuli (and our limited number of stimuli), we used only the video game paradigm in Experiments 2A and 2B, which brought the added benefit of being arguably a more ecologically valid encoding phase.

Method

Participants. Sixty-three participants from the University of Toronto, Scarborough, participated in Experiment 2A, and 30 participated in Experiment 2B. In each case, the experiment lasted for one half hour and participants earned a 0.5 bonus credit toward their introductory psychology course. All participants had normal or corrected-to-normal vision, and all were sufficiently proficient in English to complete the task without difficulty.

Procedure. The procedure was identical to that of Experiment 1B. The only difference was that instead of blurred characters, multifeature characters were used in Experiment 2A and “minimal-feature” characters were used in Experiment 2B.

Materials and Apparatus. A set of characters was created by combining one of four possible base shapes (i.e., square, circle, rectangle, flying block) with 1 of 10 possible colors, resulting in a total of 40 unique combinations of shape and color. When a base shape and color were selected, a multifeature character was created by assembling different-sized base shapes and a set of eyes into unique patterns. For example, if *square* and *purple* were selected, several different-sized purple squares were arranged to create a multifeature character. Figure 1 shows examples of such characters. Although each character had a distinctive arrangement of shapes, each was composed of shapes and colors similar to those found in other multifeature characters. Hence, we believed that this set of characters was sufficiently distinctive (although less so than normal characters) and fairly similar to one another.

To create minimal-feature characters (see Figure 1), one large base shape was used with two small base shapes to create a two-legged character. Because all were composed of similar shapes and colors, minimal-feature characters were assumed to be even more similar to one another than were multifeature characters. Likewise, the arrangements of their shapes were identical to each other—all minimal-feature characters possessed only two legs and a large base shape on top. Thus, by removing distinctive features, but leaving common features, we indirectly increased the interitem similarity of minimal-feature characters.

Results and Discussion

The results of Experiments 2A and 2B are shown in Table 1. Both experiments were analyzed separately in 2 (old vs. new) $\times 2$ (normal vs. multi-/few-/minimal-feature characters) within-subjects ANOVAs.

Experiment 2A. No overall difference was found between normal and multifeature characters [$F(1,62) < 1$]; however, participants could discriminate old items from new ones [$F(1,62) = 227.39$, $MS_e = 0.02$, $p < .01$, $\eta^2 = .79$], and a significant interaction was observed [$F(1,62) = 33.26$, $MS_e = 0.01$, $p < .01$, $\eta^2 = .35$]. Further analysis of the interaction revealed that a mirror pattern was obtained, in that participants were more likely to false alarm [$t(62) = 4.04$, $p < .01$] and less likely to hit [$t(62) = 3.07$, $p < .01$] to multifeature characters than to normal characters.

Experiment 2B. In Experiment 2B, participants were able to discriminate old from new characters [$F(1,29) = 66.65$, $MS_e = 0.02$, $p < .01$, $\eta^2 = .70$] and responded “old” more often to minimal-feature characters than to normal characters [$F(1,29) = 41.77$, $MS_e = 0.05$, $p < .01$, $\eta^2 = .59$]. A significant interaction was also observed [$F(1,29) = 42.30$, $MS_e = 0.02$, $p < .01$, $\eta^2 = .59$], which was analyzed further with paired-sample t tests. A concordant pattern was produced in that there were more hits [$t(29) = 2.14$, $p < .05$] and false alarms [$t(29) = 9.15$, $p < .01$] for minimal-feature characters than for normal characters.

Thus, in Experiments 2A and 2B, as interitem similarity increases, both hits and false alarms increase, and the mirror pattern changes to a concordant pattern. Starting with a mirror pattern in Experiment 2A, a concordant pattern emerged in Experiment 2B. Given that dual-process models assume that hit-rate portions of mirror patterns are due to the influence of recollection and that item distinctiveness is assumed to aid recollection, this framework is completely consistent with the reversal of the hit-rate portion of the mirror pattern in that distinctiveness is reduced. The increase in interitem similarity is assumed to increase item familiarity in general and, therefore, accounts for the general increase in hits and false alarms. Thus, dual-process models provide a strong overall account of the transition in the patterns observed here. In the General Discussion, we return to the issue of whether single-process accounts can readily account for these data.

EXPERIMENTS 3A AND 3B Replicating the Effects

Experiments 3A and 3B served to replicate the mirror and concordant patterns found in Experiments 2A and 2B in a traditional recognition memory paradigm. Thus, in Experiment 3A recognition for normal and multifeature characters were compared, and in Experiment 3B recognition for normal and minimal-feature characters were compared. In both experiments, participants engaged in a study phase during which items were presented automatically one at a time, after which participants were given a yes/no recognition test. The predictions for this experiment were identical to those of Experiments 2A and 2B.

Method

Participants. Fifty-nine University of Toronto, Scarborough, students participated in Experiment 3A and 20 participated in Experiment 3B for one half hour to earn a 0.5 bonus credit toward their introductory psychology course. All participants had normal or corrected-to-normal vision, and all were sufficiently proficient in English to complete the task without difficulty.

Procedure. The procedure in these experiments was identical to that of Experiment 1A.

Materials and Apparatus. In Experiment 3A, normal and multifeature characters from Experiment 2A were used. In Experiment 3B, normal and minimal-feature characters from Experiment 2B were used.

Results and Discussion

Table 1 shows comparisons of the overall hits and false alarms of Experiments 3A and 3B. Both experiments were analyzed separately in 2 (old vs. new) $\times$ 2 (normal vs. multi-/minimal-feature characters) within-subjects ANOVAs.

Experiment 3A. Participants could discriminate old from new items [$F(1,58) = 379.91$, $MS_e = 0.02$, $p < .01$, $\eta^2 = .87$]. Furthermore, although there was no overall difference between responding to normal versus multifeature characters [$F(1,58) = 2.08$, $MS_e = 0.02$, $p = .16$, $\eta^2 = .04$], character type did interact with old/new status [$F(1,58) = 22.47$, $MS_e = 0.01$, $p < .01$, $\eta^2 = .28$]. Specifically, this interaction indicated that a mirror effect was present. Paired-sample t tests confirmed that normal

characters had a higher hit rate [$t(58) = 2.11$, $p < .05$] and a lower false alarm rate [$t(58) = 3.91$, $p < .01$] than did multifeature characters. These results replicate those in Experiment 2A.

Experiment 3B. Participants could discriminate old from new items reliably [$F(1,19) = 140.53$, $MS_e = 0.01$, $p < .01$, $\eta^2 = .88$]. Old/new status also interacted with character type [$F(1,19) = 38.46$, $MS_e = 0.01$, $p < .01$, $\eta^2 = .67$]. This interaction indicated that the difference in hit rates between normal and minimal-feature characters was smaller than the difference in false alarm rates. Finally, participants reliably responded "old" more to minimal-feature characters than to normal characters [$F(1,19) = 46.08$, $MS_e = 0.04$, $p < .01$, $\eta^2 = .71$], indicating the presence of a concordant effect. Paired-sample t tests support this conclusion, in that hit rates and false alarm rates were higher for minimal-feature characters than for normal characters [$t(19) = 3.27$, $p < .01$, and $t(19) = 9.33$, $p < .01$, respectively]. These results replicate those of Experiment 2B.

We have demonstrated that, after starting from a mirror pattern, similarity enhancements can be used to obtain a concordant pattern similar to those obtained in the pseudoword effect. This result is consistent with the previously outlined dual-process account, which links the mirror effect to the pseudoword effect. However, before we claim that such an account is a good explanation of these two phenomena, we turn first to a more detailed examination of our stimuli.

EXPERIMENT 4 Subjective Measures of Distinctiveness and Familiarity

The goal of the final experiment is to obtain some subjective ratings from participants about the distinctiveness and interitem similarity of the stimuli used throughout this study. That is, the stimuli in this study were designed to have varying levels of distinctiveness and interitem similarity. Largely, we have relied on the fact that the stimuli behave in predictable ways, producing mirror and concordant patterns when expected, to argue that they indeed vary in terms of interitem similarity and distinctiveness in the manner they should. To provide further support for our claims, in Experiment 4 we asked participants to provide subjective ratings of the relative interitem similarity and distinctiveness of the stimuli used in our study. If our account so far is correct, both blurred and multifeature characters should be rated as less distinctive and more similar to one another than are normal characters. Furthermore, minimal-feature characters should be roughly equivalent to multifeature characters in terms of distinctiveness, but much higher in terms of interitem similarity. If this is true, it suggests that the only difference between these two stimuli types, and thus the reason that the mirror pattern for multifeature characters changed into a concordant pattern for minimal-feature characters, was the greater interitem similarity for minimal-feature characters.

Method

Participants. Of the 60 students from the University of Toronto, Scarborough, who participated in Experiment 4, 25 rated blurred versus normal characters, 25 rated multifeature versus normal characters, and 10 rated minimal-feature versus normal characters. Participants signed up for one half hour to earn a 0.5 bonus credit toward their introductory psychology course. All participants had normal or corrected-to-normal vision, and all were sufficiently proficient in English to complete the task without difficulty.

Procedure. Participants were placed into one of three between-subjects conditions to compare characters: blurred versus normal, multifeature versus normal, or minimal-feature versus normal. After they were grouped, the participants were shown an 80-trial study phase. On each trial, participants were shown a character for 1 sec with a 0.5-sec interstimulus interval. The purpose of this study phase was to familiarize the participants with all of the characters that they would later rate for distinctiveness and interitem similarity.

After the study phase, participants were told that they would now rate the characters they had just seen and that they would be shown either a single character or a pair of characters. Participants were instructed to rate each single character on how distinct that individual character was. Specifically, participants were told "Try to think how much that individual character stands out from ALL THE OTHER CHARACTERS, and choose a number from 1 to 9. 9 indicates that the character is very distinct and different from other characters, 1 indicates that character is not distinct at all." For character pairs, participants were instructed to "indicate how similar you think they are, choose a number from 1 to 9. 9 means you think they are very similar and 1 indicates you don't think they are similar at all."

The test phase consisted of 160 trials. On 80 of those trials, participants rated the distinctiveness of a character; on the remaining 80, participants rated the similarity of two randomly chosen characters. Characters in pairs were always from the same character type. The result is that each participant provided an average distinctiveness rating for both the normal and the other character type, as well as an average interitem similarity rating.

Materials and Apparatus. The normal, blurred, multifeature, and minimal feature characters from Experiments 1, 2, and 3 were used in this experiment.

Results and Discussion

Normal versus blurred characters. The results of the overall distinctiveness and similarity ratings can be seen in Table 2. Previously, blurred characters were found to produce a mirror pattern when compared with normal characters; thus, we expected blurred characters to be rated as less distinctive and higher in interitem similarity than normal characters. Participants rated blurred characters as less distinctive than normal characters [$t(24) = 5.22$, $p < .01$]; however, there was no difference in interitem

similarity between normal and blurred characters [$t(24) = 0.16$, $p = .87$].

The failure to find a significant difference between the interitem similarity ratings of blurred characters poses a problem for our dual-process account. One explanation is that, given that the false alarm portion of the mirror effect is assumed to rely on familiarity, and given that the false alarm portion of the mirror effect was not completely reliable in Experiment 1A when the noted outliers were not removed, it may be reasonable to suppose that blurred characters do not differ largely or reliably from normal characters in terms of familiarity and that this is why the false alarm portion of the mirror effect was not entirely consistent in Experiment 1A.

A second possibility is that the subjective interitem similarity ratings were inherently limited in accurately measuring true interitem similarity. That is, it is likely that the subjective interitem similarity ratings gathered are influenced by factors other than the interitem similarity of a stimulus set. Indeed, a quick examination of Table 2 shows that the interitem similarity ratings and the distinctiveness ratings of normal characters are not wholly reliable, since they shift slightly depending on which other characters are being evaluated. These ratings, then, by being influenced by factors other than the characteristics of the specific stimuli in question, may not be sensitive enough to bring out differences in true interitem similarity in every case. For example, if the true interitem similarity levels of two stimulus classes were fairly close but nonetheless one was greater in interitem similarity than the other, subjective interitem similarity ratings may have difficulty picking up this difference because extraneous factors would be able to shift subjective ratings and hide any true difference. This may be why a significant difference between interitem similarity ratings for normal and blurred characters was not obtained.

A final possibility is that the false alarm rate portion of the mirror pattern obtained between normal and blurred characters was not being driven by interitem similarity, but by some other factor. For example, participants may expect blurred characters to be relatively difficult to process, because they do not readily represent a distinct character or object. However, blurred characters are all based on normal characters. Thus, for some blurred characters, participants also may see the normal character on which

Table 2
Means (and Standard Deviations) for Distinctiveness
and Interitem Similarity Ratings for Different Character Types in Experiment 4

	Distinctiveness				Interitem Similarity			
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
	Normal		Blurred		Normal		Blurred	
Normal vs. blurred	5.84	1.50	3.65	1.18	3.95	1.35	3.91	1.03
	Normal		Multifeature		Normal		Multifeature	
Normal vs. multifeature	5.28	1.23	3.68	1.17	3.62	1.31	4.09	1.21
	Normal		Minimal Feature		Normal		Minimal Feature	
Normal vs. minimal feature	6.43	1.59	3.07	1.40	3.99	0.82	5.62	1.19

the blurred character was based. When participants see blurred characters later on, the blurred character may remind them of the normal character and allow them to recognize blurred characters more easily than we expected. Whittlesea and Williams (1998) suggested that processing stimuli more fluently than expected can give rise to a false sense of familiarity, and in a recognition memory context, can lead to an increase in false alarms. In terms of our dual-process account, this explanation is consistent. That is, in the introduction it was mentioned that many different factors may contribute to the familiarity of a stimuli set. We decided to focus on manipulating interitem similarity to influence familiarity, but manipulation of expected fluency is no less valid an approach to influencing familiarity. Thus, the actual fluency of blurred characters may have surpassed their expected fluency, and this may have been driving an inflated sense of familiarity with respect to blurred over normal characters.

Although the failure to find a significant difference for interitem similarity ratings between normal and blurred characters was unexpected, each of these three explanations we have suggested is consistent with our general framework. Thus, this result is not necessarily counter to our dual-process account. However, we turn now to multiple- and minimal-feature characters to provide more direct support of our account.

Normal versus multifeature characters. Multifeature characters produced a mirror effect in both Experiments 2A and 3A and so were therefore predicted by our dual-process account to show less distinctiveness and greater interitem similarity than normal characters. Participants rated multifeature characters as less distinctive [$t(24) = 6.22, p < .01$] and greater in interitem similarity [$t(24) = 2.59, p < .05$] than normal characters.

Normal versus minimal-feature characters. A full concordant pattern was revealed with minimal-feature characters in both Experiments 2B and 3B; thus, we expected interitem similarity to be higher for minimal-feature characters than for normal characters. With respect to distinctiveness, we predicted that minimal-feature characters should be nearly as distinctive as multifeature and blurred characters. That is, according to our dual-process account, enhanced interitem similarity is the only difference necessary between minimal-feature and either multifeature or blurred characters to give rise to a concordant pattern.

As predicted, minimal-feature characters were greater in interitem similarity than normal characters [$t(9) = 3.74, p < .01$]; however, minimal-feature characters were also less distinctive than normal characters [$t(9) = 4.94, p < .01$]. One key assumption of our dual-process account is that the only condition necessary for a concordant pattern to emerge is that interitem similarity is increased. A mirror pattern, on the other hand, should emerge only when distinctiveness is greater and familiarity is lesser for one stimulus class over another. Thus, the finding that minimal-feature characters were greater in interitem similarity and less in terms of distinctiveness compared with normal characters suggests that perhaps we should have found a mirror pattern, not a concordant pattern, when comparing normal and minimal-feature characters.

The dual-process account is able to account for these data, in that, if we assume that an increase in interitem similarity increases both hit and false alarm rates, whereas a decrease in distinctiveness only decreases hit rates, it is possible that a simultaneous increase in familiarity and decrease in distinctiveness should result in both hits and false alarms increasing, but with hit rates increasing more slowly. In other words, if distinctiveness for minimal-feature characters is truly less than that of multifeature characters, yet a concordant pattern was found for minimal-feature characters, then our model suggests that the change in hit rate noted between multifeature and minimal-feature characters must always be less than the change in false alarm rate noted. This is because the influences of familiarity and distinctiveness oppose one another in hit rates, but not in false alarm rates.

Looking back at our data, false alarm rates for multifeature characters in the traditional study list and the video game study phase were .44 and .35, respectively. False alarm rates increased to .65 and .63, respectively, when minimal-feature characters were used. Thus, increases of .21 and .28 was observed. Hit rates to multifeature characters in the traditional study list and the video game study phase were .68 and .55, respectively, and these increased to .80 and .68 when minimal-feature characters were used. Thus, increases of .12 and .13 were observed. Consistent with our dual-process account, hit rates were rising, as would be expected if interitem similarity were being increased, but were rising more slowly than were false alarm rates, as would be expected if distinctiveness were also decreasing.

Further support for this notion comes from an examination of the past work with the pseudoword effect. To examine the pseudoword effect in more detail, we compiled the results from several previous studies that compared recognition memory for words and pseudowords (Cleary et al., 2007; Gardiner & Java, 1990; Greene, 2004; Hockley & Nieuwadoski, 2001; Joordens et al., 2008; Rao & Proctor, 1984; Whittlesea & Williams, 1998, 2000; Wixted, 1992). For each study, we subtracted the hit or false alarm rates for the pseudowords from the hit or false alarm rates for the normal words. In all, 26 experiments were analyzed. The results are plotted in Figure 2. A positive value indicates more hits or false alarms to pseudowords than to normal words. As can be seen in Figure 2, on average the pseudoword effect in the false alarms is larger ($M = .12, SD = .11$) than in the hit rates ($M = .03, SD = .08$) [$t(50) = 3.35, p < .01$].

Critically, this finding suggests that indeed pseudowords may be less distinctive than words. What these new results mean for our account of the pseudoword effect and other concordant patterns is this: Our dual-process account suggests that the only condition necessary to give rise to a concordant pattern is that one stimuli class must be greater in terms of familiarity than another. This does not imply that the distinctiveness of the two stimulus classes must be equal. However, if the stimulus class that evokes a concordant pattern is found to have less distinctiveness than the control stimulus class, then the size of the concordant pattern in the hit rates must be less than

Figure 2. Pseudoword effect found in hits and false alarms from several experiments comparing pseudowords or rare words with words (i.e., Cleary et al., 2007; Gardiner & Java, 1990; Greene, 2004; Hockley & Nieuwadoski, 2001; Joordens et al., 2008; Rao & Proctor, 1984; Whittlesea & Williams, 1998, 2000; Wixted, 1992).

the size of the pattern in the false alarm rates, according to our account. This is because, in this case, the influence of distinctiveness will oppose the influence of familiarity in the hit rates, and thus attenuate the concordant pattern. In other words, a concordant pattern emerges when one stimulus class is more familiar than another and when that familiarity difference is enough to overcome any potential distinctiveness disadvantage.

GENERAL DISCUSSION

As researchers come to understand memory phenomena such as the frequency-based mirror effect (Glanzer & Adams, 1985) and the pseudoword effect (Greene, 2004), they often make claims about the ways in which stimulus features such as preexperimental familiarity, interitem similarity, and distinctiveness differentially affect familiarity versus recollection. The present article provides a direct test of some of these claims via a set of stimuli created to vary systematically with respect to interitem similarity and distinctiveness. In so doing, we were able to replicate mirror patterns and concordant patterns in predictable ways, and we did so in the context of what might be considered a more ecologically valid encoding context.

In the beginning of this article we outlined a dual-process account that unifies the frequency-based mirror effect and the pseudoword effect: that the mirror effect arises because high-frequency words are assumed to be less distinctive

and more familiar than are low-frequency words. On the other hand, the pseudoword effect arises because pseudowords lack semantics and therefore can be discriminated only on the basis of orthography. As a result, they have more interitem similarity. This interitem similarity gives rise to a larger sense of familiarity, which leads to more hits and more false alarms to pseudowords over words.

Our study was conducted not with words, but with characters specially designed to mimic the features hypothesized to give rise to effects like the mirror effect in words. Using both a typical old/new recognition paradigm (Experiment 1A) and a distracting video game study session (Experiment 1B), we obtained evidence for a mirror-like effect. We replicated this effect by using a radically different set of stimuli (Experiments 2A and 3A).

Furthermore, we demonstrated that, starting with a mirror pattern (Experiments 2A and 3A), data can be transformed to a concordant pattern by increasing the interitem similarity, and therefore the familiarity, of items (Experiments 2B and 3B). The results of Experiments 2B and 3B also nicely demonstrate that as discriminability is reduced, a mirror pattern is not necessarily obtained, as might be predicted from a simple single-process account, because in both experiments, there was evidence of a concordant pattern instead.

One problem with the results of Experiment 4, however, is that the blurred characters we examined, which gave rise to a mirror pattern, were found not to differ from normal characters with respect to interitem similarity. Although we offered several explanations for this result, one explanation we did not consider is that perhaps this is more consistent with a single-process account than a dual-process account. That is, if blurred characters are simply less distinctive than normal characters, we could consider them simply less memorable. If this fact alone gives rise to a mirror pattern, the results are consistent with a single-process account.

Single-process accounts of mirror and concordant patterns are more parsimonious than dual-process accounts, and so should be preferred if they can account for the same data as dual-process accounts. As mentioned before, Malmberg et al. (2004) were able to account for concordant patterns by using the single-process REM model. Results like these suggest that single-process accounts can, in principle, account for data such as ours.

However, one problem with extending a single-process account, such as REM, to our data is that increases in interitem similarity should not necessarily lead to concordant patterns, as we observed. That is, one way to implement increasing interitem similarity in the REM model could be to increase a parameter (g) that determines how common the features are that make up stimuli. The problem is that typically this parameter is manipulated to give rise to a mirror effect, not a concordant pattern.²

Experiment 4 did show that the mirror and concordant patterns in Experiments 2A–3B likely were occurring in the manner we hypothesized under our dual-process account. That is, a mirror pattern emerges when one stimulus class is both more familiar and less distinctive than another, and a concordant pattern emerges when one stim-

ulus class is generally more familiar than another. However, we found evidence that the minimal-feature characters, which gave rise to a concordant pattern, were not just more familiar, but also had less distinctiveness than expected. Yet, because our account suggests that distinctiveness can give rise to hit rates, the finding that minimal-feature characters were less distinctive than expected does not undermine our account, as a reanalysis of our previous results showed that, for minimal-feature characters, the magnitude of the hit rate advantage was smaller than that of the false alarm rate advantage. Thus, this finding is consistent with our dual-process account.

Finally, the generalizability of our results is also worth noting. Our findings suggest that a typical old/new recognition memory procedure is not required for the memory effects to be observed, and more importantly, that the effects still emerge, even in cases when cognitive demands can be quite involving. This supports the notion that such effects as the mirror and pseudoword effects generalize beyond the lab and into the real world, where items rarely are presented sequentially in a list for individuals to study and later recognize.

Although list-learning experiments will always be needed to weed out potential confounds, video game experiments may also be very useful. Because our video game experiments seemed to replicate effects produced in the lab, video game memory experiments may be useful in future research investigating the memories of populations that are difficult to test via traditional means. For example, it is conceivable that children may be more interested and more serious about a video game experiment than a traditional experiment. Conducting experiments as video games may be useful in mustering participant interest or simply in maintaining attention during longer experiments, especially with easily distracted populations.

AUTHOR NOTE

The research was supported by research grants from the Natural Sciences and Engineering Research Council of Canada and the Premier's Research Excellence Award. We thank Rey Rondina for collecting some of the pilot data. Address correspondence to J. D. Ozubko, Department of Psychology, University of Waterloo, 200 University Ave. West, Waterloo, ON, N2L 3G1 Canada (e-mail: jdozubko@uwaterloo.ca).

REFERENCES

- BALOTA, D. A., BURGESS, G. C., CORTESE, M. J., & ADAMS, D. R. (2002). The word-frequency mirror effect in young, old, and early-stage Alzheimer's disease: Evidence for two processes in episodic recognition performance. *Journal of Memory & Language*, **46**, 199-226.
- CLEARY, A. M., MORRIS, A. L., & LANGLEY, M. M. (2007). Recognition memory for novel stimuli: The structural regularity hypothesis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **33**, 379-393.
- CRISS, A. H., & SHIFFRIN, R. M. (2004). Interactions between study task, study time, and the low-frequency hit rate advantage in recognition memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 778-786.
- DEESE, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, **58**, 17-22.
- GARDINER, J. M., & JAVA, R. I. (1990). Recollective experience in word and nonword recognition. *Memory & Cognition*, **18**, 23-30.
- GARDINER, J. M., & JAVA, R. I. (1993). Recognition memory and awareness: An experiential approach. *European Journal of Cognitive Psychology*, **5**, 337-346.
- GERACI, L., & RAJARAM, S. (2002). The orthographic distinctiveness effect on direct and indirect tests of memory: Delineating the awareness and processing requirements. *Journal of Memory & Language*, **47**, 273-291.
- GILLUND, G., & SHIFFRIN, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, **91**, 1-67.
- GLANZER, M., & ADAMS, J. K. (1985). The mirror effect in recognition memory. *Memory & Cognition*, **13**, 8-20.
- GLANZER, M., & ADAMS, J. K. (1990). The mirror effect in recognition memory: Data and theory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 5-16.
- GREENE, R. L. (2004). Recognition memory for pseudowords. *Journal of Memory & Language*, **50**, 259-267.
- HIGHAM, P. A., & VOKEY, J. R. (2000). Judgment heuristics and recognition memory: Prime identification and target-processing fluency. *Memory & Cognition*, **28**, 574-584.
- HILFORD, A., GLANZER, M., & KIM, K. (1997). Encoding, repetition, and the mirror effect in recognition memory: Symmetry in motion. *Memory & Cognition*, **25**, 593-605.
- HINTZMAN, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace model. *Psychological Review*, **95**, 528-551.
- HINTZMAN, D. L., & CURRAN, T. (1997). Comparing retrieval dynamics in recognition memory and lexical decision. *Journal of Experimental Psychology: General*, **126**, 228-247.
- HIRSHMAN, E., FISHER, J., HENTHORN, T., ARNDT, J., & PASSANANTE, A. (2002). Midazolam amnesia and dual-process models of the word-frequency mirror effect. *Journal of Memory & Language*, **47**, 499-516.
- HOCKLEY, W. E., & NIEWIADOMSKI, M. W. (2001). Interrupting recognition memory: Tests of a criterion-change account of the revelation effect. *Memory & Cognition*, **29**, 1176-1184.
- JACOBY, L. L., & DALLAS, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, **110**, 306-340.
- JACOBY, L. L., & WHITEHOUSE, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, **118**, 126-135.
- JOORDENS, S., & HOCKLEY, W. E. (2000). Recollection and familiarity through the looking glass: When old does not mirror new. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 1534-1555.
- JOORDENS, S., & MERIKLE, P. M. (1992). False recognition and perception without awareness. *Memory & Cognition*, **20**, 151-159.
- JOORDENS, S., OZUBKO, J. D., & NIEWIADOMSKI, M. W. (2008). Featuring old/new recognition: The two faces of the pseudoword effect. *Journal of Memory & Language*, **58**, 380-392.
- MALMBERG, K. J., ZEELLENBERG, R., & SHIFFRIN, R. M. (2004). Turning up the noise or turning down the volume? On the nature of the impairment of episodic recognition memory by midazolam. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 540-549.
- MURDOCK, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, **89**, 609-626.
- RAO, K. V., & PROCTOR, R. W. (1984). Study-phase processing and the word frequency effect in recognition memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **10**, 386-394.
- REDER, L. M., NHOUYVANISVONG, A., SCHUNN, C. D., AYERS, M. S., ANGSTADT, P., & HIRAKI, K. (2000). A mechanistic account of the mirror effect for word frequency: A computational model of remember-know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 294-320.
- ROEDIGER, H. L., III, & McDERMOTT, K. B. (1995). Creating false memories: Remembering words not presented on lists. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 803-814.
- ROTELLO, C. M., & MACMILLAN, N. A. (2006). Remember-know models as decision strategies in two experimental paradigms. *Journal of Memory & Language*, **55**, 479-494.
- SCHMIDT, S. R. (1991). Can we have a distinctive theory of memory? *Memory & Cognition*, **19**, 523-542.
- SHIFFRIN, R. M., HUBER, D. E., & MARINELLI, K. (1995). Effects of

- category length and strength on familiarity in recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 267-287.
- SHIFFRIN, R. M., & STEYVERS, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, **4**, 145-166.
- SOMMERS, M. S., & LEWIS, B. P. (1999). Who really lives next door: Creating false memories with phonological neighbors. *Journal of Memory & Language*, **40**, 83-108.
- VALENTINE, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology*, **43A**, 161-204.
- WHITTLESEA, B. W. A., & WILLIAMS, L. D. (1998). Why do strangers feel familiar, but friends don't? A discrepancy-attribution account of feelings of familiarity. *Acta Psychologica*, **98**, 141-165.
- WHITTLESEA, B. W. A., & WILLIAMS, L. D. (2000). The source of feelings of familiarity: The discrepancy-attribution hypothesis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 547-565.
- WICKELGREN, W. A. (1966). Numerical relations, similarity, and short-term recognition memory for pairs of digits. *British Journal of Psychology*, **57**, 263-274.
- WIXTED, J. T. (1992). Subjective memorability and the mirror effect. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 681-690.
- WIXTED, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, **114**, 152-176.
- YONELINAS, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory & Language*, **46**, 441-517.
- ZAKI, S. R., & NOSOFSKY, R. M. (2001). Exemplar accounts of blending and distinctiveness effects in perceptual old-new recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 1022-1041.

NOTES

1. Although the distinction between similarity and distinctiveness sometimes can be difficult to make, especially because increasing similarity seems to, by necessity, decrease distinctiveness, we feel it is still possible to discuss these concepts separately. For example, to remove distinctive features, we blurred all characters in a similar manner. It would have been possible to remove distinctive features without blurring characters, in which case characters would not necessarily become more similar to one another. That is, characters would lose distinctiveness, but not necessarily come to resemble each other any more than usual if we had, for example, simply removed all of the characters' eyes and arms. Thus, although our blurring manipulation had opposite-acting effects on distinctiveness and similarity, it should be noted that this was intentional and that distinctiveness and similarity are not necessarily the same concept, although they definitely are related to one another.

2. We thank Jason Arndt for making this point to us during the review process.

(Manuscript received September 19, 2006;
revision accepted for publication July 9, 2008.)