

Research Article
Psychological Science

Conceptual Penetration of Visual Processing

¹Institute for Research in Cognitive Science, ²Center for Cognitive Neuroscience, and
³Department of Psychology, University of Pennsylvania

Corresponding Author:

Gary Lupyan, Institute for Research in Cognitive Science, University of Pennsylvania, 3401
Walnut St., Philadelphia, PA 19104
E-mail: lupyan@sas.upenn.edu

Conceptual Penetration of Visual Processing

Gary Lupyan^{1,2}, Sharon L. Thompson-Schill^{1,2,3}, and Daniel Swingley^{1,3}

Abstract

In traditional hierarchical models of information processing, visual representations feed into conceptual systems, but conceptual categories do not exert an influence on visual processing. We provide evidence, across four experiments, that conceptual information can in fact penetrate early visual processing, rather than merely biasing the output of perceptual systems. Participants performed physical-identity judgments on visually equidistant pairs of letter stimuli that were either in the same conceptual category (*Bb*) or in different categories (*Bp*). In the case of nonidentical letters, response times were longer when the stimuli were from the same conceptual category, but only when the letters were presented sequentially. The differences in effect size between simultaneous and sequential trials rules out a decision-level account. An additional experiment using animal silhouettes replicated the major effects found with letters. Thus, performance on an explicitly visual task was influenced by conceptual categories. This effect depended on processing time, immediately preceding experience, and stimulus typicality, which suggests that it was produced by the direct influence of category knowledge on perception, rather than by a postperceptual decision bias.

Keywords

visual processing, categorization, modularity of vision, top-down effects, perceptual learning, knowledge-based vision

Received 6/29/09; Revision accepted 9/1/09

Extracting meaning from a perceptual signal is an act of categorization, requiring the observer to highlight some stimulus features while overlooking (abstracting over) others that are irrelevant to a given category (Harnad, 2005). Humans categorize objects with incredible speed. By some measures, basic-level categorization occurs in parallel with object detection; by the time an object is detected, it is already categorized to some degree (Grill-Spector & Kanwisher, 2005; cf. Mack, Gauthier, Sadr, & Palmeri, 2008). When participants are monitoring for a known object category, neural signatures of categorization can be measured as early as 150 ms after stimulus onset (VanRullen & Thorpe, 2001; cf. Johnson & Olshausen, 2003). It has been argued that such rapid categorization can be achieved by biasing visual processing to favor low-level features diagnostic of the relevant category (Delorme, Rousselet, Mace, & Fabre-Thorpe, 2004; Johnson & Olshausen, 2003; McCotter, Gosselin, Maccabee, & Schyns, 2005; Schyns & Oliva, 1999).

In traditional hierarchical models of information processing, early visual processing feeds into conceptual systems, but the latter do not influence visual processes (Glezer, Jiang, & Riesenhuber, 2009; Pylyshyn, 1999; Riesenhuber & Poggio, 2000). Here, we provide evidence, from four experiments, that conceptual information can in fact penetrate early visual processing, rather than merely biasing the output of perceptual systems.

There is often a conflict between the physical and conceptual relationships between objects. Consider the letter pairs *B-b* and *B-p*. The physical relationship between the letters is identical in the two pairs—*B* is physically equidistant from *p* and *b*. However, conceptually, *B* is more similar to *b* than to *p*. A conceptual effect on perception—greater perceived similarity between *B* and *b* than between *B* and *p*—can arise in different ways.

One possibility, consistent with studies of perceptual learning (Goldstone, 1994, 1998; Goldstone, Steyvers, Spencer-Smith, & Kersten, 2004; Harnad, 1987; Kuhl, 1994; Livingston, Andrews, & Harnad, 1998; Notman, Sowden, & Ozgen, 2005; Schyns & Rodet, 1997), is that long-term experience categorizing *B* and *b* as members of the same category increases their perceptual similarity, and, conversely, that categorizing *B* and *p* into different categories decreases their similarity, through a gradual, bottom-up retuning of visual feature detectors. Empirical demonstrations of these effects typically involve perceptual comparisons of novel items before and after categorization experience.

However, changes in similarity structure can also result from on-line, transient modulation of perceptual processing by higher-level conceptual representations (a top-down process). Consistent with this view are studies showing that visual categorical perception can sometimes be disrupted by verbal interference (which arguably acts to block conceptual effects on perception; A. Gilbert, Regier, Kay, & Ivry, 2006; Winawer et al., 2007). Such a modulatory account is compatible with electrophysiological findings that show very rapid top-down modulation of early visual processing (e.g., V1) starting 10 to 50 ms after stimulus onset (Hupe et al., 2001; for reviews, see C. Gilbert & Sigman, 2007; Lamme & Roelfsema, 2000). Such rapid modulation is made possible by a surprisingly fast posterior-to-anterior activation flow, with visual cortex activating prefrontal cortex in as little as 30 ms (Foxy & Simpson, 2002). Foxy and Simpson argued that these results “provide a context for appreciating the 100–400 ms of processing necessary prior to response output in humans, [which leaves] ample time for multiple cortical interactions at all levels of the system” (p. 145).

Finally, some researchers have argued that apparent conceptual effects on perception do not involve alterations of visual processing at all, and can be fully explained by high-level decision biases (e.g., Pylyshyn, 1999).

In this article, we present four experiments aimed at testing the idea that conceptual effects on perceptual processing are best described in terms of a dynamic top-down process. We hypothesized that the perceived visual similarity between two stimuli will become increasingly affected by the category membership of the stimuli over the course of several hundred milliseconds, as the conceptual category information feeds back on lower-level perceptual representations.

Experiment 1

A highly reliable measure of visual similarity is the time it takes to determine whether two stimuli are physically different. Increases in perceived visual similarity have a behavioral signature of greater response times (RTs). Thus, effects of conceptual categories on perceptual processing can be studied by manipulating the conceptual relationship between stimuli while keeping the physical relationship constant. We define the *category effect* as the difference in RTs between within-conceptual-category pairs (*Bb*) and between-category pairs (*Bp*). We hypothesized that when the two stimuli were presented simultaneously, participants would respond before the conceptual representations exerted a significant influence on the ongoing visual processing. Thus, we predicted a small or nonexistent category effect for simultaneously presented stimuli. However, with greater time for processing and categorizing one of the items, its visual representation would become more similar to the representations of other members of the same category or more different from the representations of members of different categories.

We thus contrasted the category effect during simultaneous judgments and sequential judgments. Our prediction was that it would take longer to respond “different” to within-category trials (*Bb*) than to between-category trials (*Bp*), and that the size of this category effect would be larger for sequentially presented than for simultaneously presented pairs.

Method

Twelve University of Pennsylvania undergraduates participated in Experiment 1 for course credit. They were asked to perform a speeded same/different task, responding via a keyboard button press to indicate “same” if the letters in a pair were physically identical, and “different” otherwise. All pairs were composed from the set of the letters *B*, *b*, and *p*. The critical pairs were the two *different* pairs: *Bb* (within category) and *Bp* (between category). Pixel by pixel, the letters *B* and *b* were as similar to each other as the letters *B* and *p*. Any differences in responding to these two pairs could not be due to differences in visual similarity, and thus would have to be attributed to conceptual differences (Lupyan, 2008a).

On half of the trials, the letters were presented simultaneously. On the remaining half, the first letter was displayed on the screen for 150, 300, 450, or 600 ms before the second letter appeared. The first letter remained on the screen during the variable delay (Fig. 1).¹ To discourage the allocation of attention to a specific part of the display during the delay, we varied the location of the letters; on each trial, the two stimuli ($\sim 0.7^\circ \times 0.9^\circ$) appeared in two of four possible (randomly selected) locations equidistant ($\sim 1.5^\circ$) from a central fixation cross. There were no memory demands, as both stimuli remained on the screen until a response was made. The intertrial interval was 750 ms.

Participants completed 12 practice trials followed by 576 experimental trials. There were equal numbers of *same* and *different* trials, and of within- and between-category trials.

Participants were encouraged to respond as quickly as possible without compromising accuracy.

A buzz sounded after incorrect responses. Hand-to-response assignment was counterbalanced across subjects.

Results and discussion

Errors were rare (< 5%). We analyzed RTs from the onset of the second stimulus. Trials responded to incorrectly were excluded from analysis, as were trials with RTs shorter than 200 ms or longer than 1,200 ms (3.3%). Hand-to-response assignment did not interact with any variables of interest ($F_s < 1$).

RTs were analyzed using repeated measures analyses of variance (ANOVAs).

Participants responded more quickly to *same* trials ($M = 550$ ms) than to *different* trials ($M = 577$ ms), $F(1, 11) = 10.16, p = .009$ —an example of a previously observed “fast-same” effect (Posner, 1978). Unsurprisingly, the extra processing time available in sequential trials yielded faster RTs than were observed in simultaneous trials, $F(1, 11) = 26.70, p < .0005$.

The critical prediction concerned the *different* trials. We predicted that RTs on these trials would be faster when the two stimuli were categorically different (Bp) than when they were categorically identical (Bb), but only when the first stimulus received extra processing time (i.e., when there was a nonzero stimulus onset asynchrony, or SOA). RTs were significantly slower on within-category trials ($M = 589$ ms) than on between-category trials ($M = 565$ ms), $F(1, 11) = 9.91, p = .009$. Critically, there was a significant interaction between category type (same vs. different) and SOA (zero vs. nonzero), $F(1, 11) = 7.99, p = .016$ (Fig. 2a). When the two letters

were presented simultaneously, there was an unreliable 11-ms difference between the two trial types, $t(11) = 1.78, p > .1$ (two-tailed). In contrast, the 38-ms difference observed on the sequential trials was significant, $t(11) = 3.32, p = .007$. There was no evidence of speed-accuracy trade-offs. Participants were not only faster, but also slightly (1%) more accurate on the between-category trials, $F(1, 11) = 2.98, p = .11$. There were no effects of presentation order (e.g., Bp vs. pB), $t < 1$.

The magnitude of the category effect did not vary reliably among the nonzero SOAs, but as Figure 2a illustrates, the effect became more reliable as SOA increased. Simultaneous Levine tests showed that variance was significantly higher at the 150-ms SOA than at the other delays and decreased from the 300-ms to the 450-ms SOA, $ps < .01$.

It may at first appear that these results are compatible with two accounts: a conceptual-penetration account, in which categories affect visual representations, and a decision-level account, in which same-category pairs (e.g., Bb) introduce a "same" response bias that competes with the desired "different" response. According to the latter account, the decision to respond "same" or "different" is a linear combination of two input channels, one used to make the decision on the basis of visual features and one used to make the decision on the basis of conceptual features (Hawkins et al., 1990); the existence (and size) of the category effect would then depend on the weights assigned to the channels. However, in the present paradigm, neither the visual nor the conceptual response alternatives can compete until both stimuli are visible (i.e., the channel weights are effectively 0 until both stimuli are present). Thus, a decision-level account predicts a category effect that is independent of SOA (see A Test of the Decision-Level Model Account in the Supplemental Material available online). In contrast, according to the conceptual-penetration account, categorization of the first stimulus influences the perceptual

processing of the second, and the feedback effect of categories on perceptual processing is more limited when the stimuli are presented simultaneously or with a short SOA.

Experiment 2

We hypothesized that conceptual effects on perceptual processing are modulated by the strength of the association between the visual form and the conceptual category (Lupyan, 2007, 2008a). Manipulations that disrupt this association are therefore predicted to reduce or eliminate the conceptual-penetration effect (the RT difference between *Bb* and *Bp* trials). In Experiment 2, we achieved such disruption by repeating the procedures of Experiment 1 with the stimuli rotated 90°. Rotating the stimuli preserved the physical relationship between them, but we assumed it would partially disrupt the association between the visual forms and the conceptual categories. Fifteen University of Pennsylvania undergraduates participated for course credit.

Trials responded to incorrectly were excluded from analysis, as were trials with RTs shorter than 200 ms or longer than 1,200 ms (1.0%). In contrast to Experiment 1, which showed a large effect of category type on sequential trials, Experiment 2 showed no evidence of a category effect, as confirmed by a repeated measures ANOVA with category effect as a fixed factor $F(1, 14) < 1$ (Fig. 2b). Accuracy was high ($M = 96.5\%$) and did not significantly differ between conditions; the direction of accuracy differences mirrored the direction of RT analyses.

To compare Experiments 1 and 2 directly, we tested the three-way Experiment \times SOA (zero vs. nonzero) \times Category Type interaction, which proved to be significant, $F(1, 25) = 4.03$ $p = .048$. Further analysis revealed that the Experiment \times Category Type interaction was significant for the sequential trials, $F(1, 25) = 7.10$, $p = .010$, but not for the simultaneous trials, $F < 1$. That is, although responses to the sequential between-category (*Bp*) trials were

significantly faster than responses to the sequential within-category (*Bb*) trials in Experiment 1 ($M_{\text{between}} = 535$ ms, $M_{\text{within}} = 573$ ms), this pattern did not hold for sequential trials in Experiment 2 ($M_{\text{between}} = 577$ ms, $M_{\text{within}} = 580$ ms).

These results confirm the prediction that manipulations that maintain all the low-level visual components of the stimuli, but disrupt their relation to the conceptual category, are sufficient to disrupt conceptual effects on visual processing.

Experiment 3

We reasoned that just as disrupting the association between the visual form and its category decreased the influence of the conceptual category on visual processing, strengthening the association would increase the strength or speed of the category effect.

Method

This experiment was identical to Experiment 1 except that the main task was preceded by a 5-min overt categorization task in which participants were instructed to respond with the “b” key to *Bs* and *bs* and with the “p” key to *ps*. On each trial, one letter appeared by itself randomly in one of the four positions used in Experiment 1 and remained on the screen until a response. There were 80 *b* trials and 80 *p* trials randomly intermixed. Following this categorization task, the same/different judgment trials were presented as in Experiment 1. Nineteen University of Pennsylvania undergraduates participated for course credit.

Results and discussion

Mean RT for the overt classification phase was 492 ms (decreasing from 532 ms for the first 40 trials to 472 ms for the last 40). Classification errors were rare (< 5%). As in Experiment 1, the

RT analysis of the same/different task excluded incorrect responses and trials with RTs shorter than 200 ms or longer than 1,200 ms (2.2%).

A repeated measures ANOVA conducted on the *different* trials in the main task revealed significant effects of SOA (zero vs. nonzero), $F(1, 18) = 72.89, p < .0005$, and category type (within- vs. between-category trials), $F(1, 18) = 16.55, p = .001$. Unlike in Experiment 1, there was no interaction between these factors, $F < 1$ (Fig. 3).

Planned t tests (two-tailed) revealed that, unlike in Experiment 1, there was a significant category effect when the stimuli were presented simultaneously ($M = 18$ ms), $t(18) = 4.62, p < .0005$. The effect at the 150-ms SOA was of comparable magnitude to that observed in Experiment 1 (~ 31 ms), but in this case was reliably greater than 0, $t(18) = 2.43, p = .026$. Unlike in Experiment 1, there was a reliable main effect of category type on accuracy. Accuracy was significantly greater for the between-category trials ($M = 96.3\%$) than for the within-category trials ($M = 93.2\%$), $F(1, 18) = 19.71, p < .0005$. There was no interaction between category type and SOA, $F < 1$. The effect of category type on accuracy was greater in Experiment 3 than in Experiment 1, $F(1, 29) = 4.29, p = .039$.

In an unexpected departure from Experiment 1, RTs for within-category trials were no greater than RTs for between-category trials for SOAs of 450 and 600 ms. For these longer SOAs, the category effect was significantly greater in Experiment 1 than in Experiment 3, $F(1, 29) = 9.46, p = .005$. This result provides an important constraint on the proposed account, suggesting that the category effect not only has a dynamic onset, but also has an offset.

Strengthening the association between the visual form and the conceptual category appears to shift the timing of the category effect. We tested this prediction by identifying the time point at which the category effect was largest for each subject and comparing the means of the Gaussian

fits between Experiments 1 and 3 (Fig. 4). A two-tailed t test showed the peak effect timing to be significantly earlier for Experiment 3 (245 ms) than for Experiment 1 (375 ms), $t(29) = 2.19$, $p = .036$.

In addition to comparing the timing of the category effect between experiments, we examined whether subjects who were faster in categorizing the stimuli during the pretask categorization trials demonstrated an earlier category effect in the main task. Categorization RTs did not predict the overall size of the category effect, $r = -.05$, but did predict the category effect for simultaneous judgments, $r = -.755$, $p < .001$. The most parsimonious account of these results is that faster categorization induces an earlier, but not necessarily larger, category effect.

To summarize, a brief “training” session in which participants categorized (already highly overlearned) letter stimuli caused category knowledge to affect simple same/different judgments more quickly, such that the effect was observed even at SOAs of 0. Moreover, the fastest categorizers were the most likely to show a reliable category effect at an SOA of 0. Category training also led to reduced accuracy on within-category trials relative to between-category trials.

Rather than changing the overall magnitude of the category effect, classification training appeared to shift the category effect earlier in time. Although a decision-level account also predicts that classification training would produce a more robust category effect (insofar as the decision based on category identity would compete more effectively with a decision based on physical identity), a decision-level account does not predict that categorization would shift the onset and offset of the effect, because, as mentioned earlier, category-level competition cannot begin until both stimuli are visible.

Experiment 4

The purpose of Experiment 4 was to generalize the conceptual-penetration effect to a richer set of stimuli. These richer stimuli allowed us to introduce greater perceptual variability and examine effects of stimulus typicality. Insofar as more typical stimuli are categorized more quickly and more reliably, processing a typical stimulus should lead to a stronger category effect than processing an atypical stimulus.

Method

Twenty-one subjects participated in Experiment 4 for course credit. One was excluded because of a high error rate (2.77 *SD* above the mean). The design was identical to that of Experiment 1 except that the letters were replaced by richer stimuli: silhouettes of cats and dogs (Quinn, Eimas, & Tarr, 2001; see Fig. 5). These images have the advantage of being meaningful and easily classified and of forcing participants to rely on global shape information rather than local image properties. To further increase variability, we manipulated the orientation of the silhouettes from trial to trial, although on any given trial, both animals faced to the right or to the left. Each participant completed 408 trials. All stimulus pairs were balanced in a 2 (same vs. different) \times 2 (within vs. between category) \times 2 (simultaneous vs. sequential) design. A separate group of 6 participants performed 144 trials of speeded classification (analogous to the first part of Experiment 3). These data were used to classify the stimuli into typical and atypical exemplars, in order to predict the size of the category effect for different combinations of stimuli.

Results and discussion

Errors were rare (< 5%). Trials responded to incorrectly were excluded from analysis, as were trials with RTs shorter than 200 ms or longer than 1,500 ms (3.9%). RTs were faster for *same* trials ($M = 727$ ms) than for *different* trials ($M = 763$ ms), $F(1, 19) = 11.27$, $p = .003$. For the

different trials, participants were always slower to respond to cat_a-cat_b trials ($M = 814$ ms) than to cat-dog trials ($M = 747$ ms), $F(1, 19) = 61.36, p < .0005$. This difference remained relatively constant for all SOAs, which suggests that the cats were more physically similar to each other than to dogs (i.e., conceptual similarity was confounded with physical similarity).

In contrast, the physical similarity between two dogs was not greater than the physical similarity between cats and dogs. Participants were as fast to respond “different” on dog_a-dog_b trials as on dog-cat trials, but only when the two pictures appeared simultaneously (Fig. 6). When stimuli were presented sequentially, RTs for between-category trials became faster than RTs for within-category trials, resulting in a significant Category Type \times SOA interaction, as in Experiment 1, $F(1, 19) = 6.59, p = .019$. Planned comparisons did not reveal a significant category effect for simultaneous trials, $F < 1$, and revealed a marginal effect for sequential trials, $F(1, 19) = 3.95, p = .061$. As is evident in Figure 6, the effect was largest for the 300-ms SOA ($M_{\text{within-category}} = 745$ ms, $M_{\text{between-category}} = 695$ ms), at which point the difference was significantly different from 0, $F(1, 19) = 4.70, p = .043$. As in Experiment 3, the effect declined for longer SOAs.

The larger diversity of stimuli in this experiment allowed us to examine effects of category strength more directly than in the previous experiments. A more rapidly categorized stimulus is predicted to influence subsequent perceptual processing more than a more slowly categorized stimulus. Therefore, we predicted that seeing a more typical category member as the first stimulus in a sequential trial would produce a greater category effect than seeing a less typical stimulus first. To test this prediction, we used the classification RTs from the 6 additional subjects to select the cats and dogs with the fastest (Fig. 5: C₁, D₁) and slowest (Fig. 5: C₃, D₃) classification RTs. We then contrasted two types of within-category *different* trials: typical \rightarrow

atypical trials ($C_1 \rightarrow C_3$ or $D_1 \rightarrow D_3$) and atypical \rightarrow typical trials ($C_3 \rightarrow C_1$ or $D_3 \rightarrow D_1$). As predicted, typical \rightarrow atypical trials yielded significantly longer RTs ($M = 797$ ms) than atypical \rightarrow typical trials ($M = 711$ ms), $F(1, 19) = 10.49, p = .004$. Recall that the stimuli being judged in the two cases were exactly the same.

Finally, we examined typicality effects as a function of presentation time. Order of presentation is undefined for simultaneous trials, but we could compare the size of the category effect for simultaneous trials and sequential trials, the latter being divided into those that presented a typical cat or dog as the first stimulus (*typical first*) and those that did not (*atypical first*). There was a significant Typicality Condition \times Category Type interaction: The category effect was larger for the typical-first than for the atypical-first trials, $F(1, 19) = 5.29, p = .03$. Moreover, the category effect increased from simultaneous to sequential trials for the typical-first trials, $t(19) = 2.94, p = .008$, but not the atypical-first trials, $t(19) < 1$.

Experiment 4 extended the main results of the earlier studies to a more diverse set of stimuli. With natural categories it is normal for within-category physical similarity to be greater than between-category physical similarity, and this was the case with cat_a - cat_b versus cat - dog pairs. For these pairs, the magnitude of the category effect did not change with SOA and can be wholly explained through differences in physical similarity (Lupyan, 2008a). In contrast, RTs to dog_a - dog_b pairs were no different from RTs to cat - dog pairs in the case of simultaneous presentations, but diverged for longer SOAs. The degree of this divergence was mediated by relative classification difficulty (measured by data collected from a separate group of participants). The category effect was larger when the more typical (easier to classify) stimulus was presented first. These results support the conclusion that category knowledge affects even very simple

perceptual judgments, and that factors such as typicality and task timing affect this process in predictable ways (Lupyan, 2007, 2008b).

General Discussion

The time subjects took to make a simple visual decision about two familiar pictures was influenced by the categorical relationship between the pictures. Judging two pictures to be physically “different” took longer when they were in the same category than when they belonged to different categories, even though the pictures were visually equidistant in the two cases (Experiments 1 and 3). This category effect disappeared when the pictures were made less meaningful through a 90° rotation (Experiment 2). Critically, the category effect was absent when the two pictures were presented simultaneously, and emerged over the course of a delay between the presentation of the first and second stimulus (Experiments 1 and 4). Providing a brief categorization session prior to the main task shifted the onset of the category effect, producing a reliable category effect even for simultaneously presented stimuli (Experiment 3). We also found that subjects who were the fastest categorizers showed the most reliable category effect on simultaneous trials. Experiment 4 extended these findings in two important ways. First, we generalized the results to a new stimulus set. Second, we showed that the size of the category effect varies as a function of typicality: The time to judge two stimuli from the same category as different is longer when the typical stimulus is presented first.

These results provide behavioral evidence of conceptual penetration of visual processing, supporting the notion of categorical perception as a dynamic process, arising from a modulation of visual representations by higher-level conceptual representations (an account consistent with the findings of A. Gilbert et al., 2006; A. Gilbert, Regier, Kay, & Ivry, 2008; Winawer et al.,

2007; see McMurray, Aslin, Tanenhaus, Spivey, & Subik, 2008, for an extension to the speech domain). Unlike demonstrations of task-demands on visual processing (e.g., Schyns & Oliva, 1999), the present results show that visual processing in the service of the very same task (physical same/different judgments) is affected by non-visual properties—the conceptual relationship between the stimuli that are being evaluated.

Although the present studies cannot establish the neural locus of the category effect, the results rule out a simple decision-level model according to which conceptual categories influence a global decision process without influencing visual representations (Mitterer, Horschig, Musseler, & Majid, 2009; Norris, McQueen, & Cutler, 2000). This model is ruled out because the visual and conceptual channels can compete only after both stimuli are visible, which means that SOA should not affect the size of the category effect. Such a model also does not predict the typicality asymmetries in Experiment 4.

Consider the finding that participants take longer to respond “different” to two dogs when a typical dog is followed by an atypical dog than when the order of the same stimuli is reversed (Experiment 4). Clearly, a response requires processing of both stimuli. According to standard accounts, the representations of the two stimuli are compared, and the speed of the decision depends on the distance between the two stimuli in a neural state space (which might reflect visual or conceptual dimensions, or both). If the representations are conceived of as static locations in state space, then, counter to the present results, there should be no asymmetry between typical-first and atypical-first trials (, i.e., $\text{distance}[\text{typical}, \text{atypical}] = \text{distance}[\text{atypical}, \text{typical}]$). In contrast, according to the conceptual-penetration account, seeing a typical dog activates a conceptual category that alters the visual processing of other dogs. This resulting increase in similarity between the first and second dogs produces a longer latency for a

“different” response. Seeing the atypical dog first invokes this same process, but to a lesser degree. (An informative discussion and visualization can be found in Spivey, 2008, pp. 273–274.) This effect bears some similarity to asymmetry effects in reasoning (i.e., reasoning from a more dominant, or typical, to a less dominant concept; Tversky, 1977).

The present results provide support to a dynamic and thoroughly interactive view of cognition (e.g., Spivey, 2008). In this view, low-level visual representations are under a constant influence of higher-level representations, and this influence leads to activity that is context and task dependent (C. Gilbert & Sigman, 2007; Kveraga, Ghuman, & Bar, 2007; Lamme & Roelfsema, 2000).

The emergent nature of the category effect may also be thought of in terms of retrieving exemplars. Although the stimuli presented to the participants (B_p and B_p) were matched in visual similarity, one can imagine that in categorizing a B , multiple exemplars, including lowercase instances, are activated, rendering the difference between a b and $B_{\text{with-activated-exemplars}}$ smaller than the difference between a b and $B_{\text{without-activated-exemplars}}$ (Cohen & Nosofsky, 2000). Regardless of whether one views the category effect as arising from the activation of multiple specific exemplars, from the activation of a category prototype, or from neural state-space dynamics that blend prototype and exemplar properties (a view embodied in connectionist accounts, e.g., Hinton & Shallice, 1991), the implications of our results are the same: Performance on an explicitly visual task (physical same/different judgments) is influenced by conceptual categories. This effect depends on processing time, immediately preceding experience, and stimulus typicality. Together, these results constitute evidence for conceptual penetration of visual processing and open the door to examinations of the neural mechanisms by which visual processing is guided by higher-level conceptual representations.

Acknowledgments

We thank Emily McDowell, Song-I Yang, and Martekuor Doodoo for their help with data collection.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interests with respect to their authorship and/or the publication of this article.

Funding

This research was supported by a National Science Foundation training grant to G.L. and by National Institutes of Health Grants R01-DC009209 and R01-MH67008 to S.L.T-S and R01-HD049681 to D.S.

Supplemental Material

Additional supporting information may be found at <http://pss.sagepub.com/content/by/supplemental-data>

Note

1. We did not include *bp* trials because they lacked a visual control; *pp* trials were excluded to ensure a balance of *same* and *different* trials. Although this meant that *b* stimuli were seen more often than *p* stimuli, this imbalance did not confound the effect because the critical comparison (*Bp* and *Bb*) trials were seen equally often. Faster responses to *b* stimuli due to practice would work against the predicted effect.

References

- Cohen, A., & Nosofsky, R. (2000). An exemplar-retrieval model of speeded same-different judgments. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 1549–1569.
- Delorme, A., Rousselet, G., Mace, M., & Fabre-Thorpe, M. (2004). Interaction of top-down and bottom-up processing in the fast visual analysis of natural scenes. *Cognitive Brain Research*, *19*, 103–113.
- Foxe, J., & Simpson, G. (2002). Flow of activation from V1 to frontal cortex in humans: A framework for defining "early" visual processing. *Experimental Brain Research*, *142*, 139–150.
- Gilbert, A., Regier, T., Kay, P., & Ivry, R. (2006). Whorf hypothesis is supported in the right visual field but not the left. *Proceedings of the National Academy of Sciences, USA*, *103*, 489–494.
- Gilbert, A., Regier, T., Kay, P., & Ivry, R. (2008). Support for lateralization of the Whorfian effect beyond the realm of color discrimination. *Brain and Language*, *105*, 91–98.
- Gilbert, C., & Sigman, M. (2007). Brain states: Top-down influences in sensory processing. *Neuron*, *54*, 677–696.
- Glezer, L.S., Jiang, X., & Riesenhuber, M. (2009). Evidence for highly selective neuronal tuning to whole words in the "visual word form area." *Neuron*, *62*, 199–204.
- Goldstone, R. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*, 178–200.
- Goldstone, R. (1998). Perceptual learning. *Annual Review of Psychology*, *49*, 585–612.
- Goldstone, R., Steyvers, M., Spencer-Smith, J., & Kersten, A. (2004). Interactions between perceptual and conceptual learning. In E. Diettrich & A. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines* (pp. 191–228). Mahwah, NJ: Erlbaum.

- Grill-Spector, K., & Kanwisher, N. (2005). Visual recognition: As soon as you know it is there, you know what it is. *Psychological Science, 16*, 152–160.
- Harnad, S. (1987). *Categorical perception: The groundwork of cognition*. New York: Cambridge University Press.
- Harnad, S. (2005). Cognition is categorization. In H. Cohen & C. Lefebvre (Eds.), *Handbook of categorization* (pp. 20–45). Amsterdam, Elsevier.
- Hawkins, H., Hillyard, S., Luck, S., Downing, C., Mouloua, M., & Woodward, D. (1990). Visual-attention modulates signal detectability. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 802–811.
- Hinton, G.E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review, 98*, 74–95.
- Hupe, J., James, A., Girard, P., Lomber, S., Payne, B., & Bullier, J. (2001). Feedback connections act on the early part of the responses in monkey visual cortex. *Journal of Neurophysiology, 85*, 134–145.
- Johnson, J.S., & Olshausen, B.A. (2003). Timecourse of neural signatures of object recognition. *Journal of Vision, 3*, 499–512. Retrieved from <http://www.journalofvision.org/3/7/4/>
- Kuhl, P. (1994). Learning and representation in speech and language. *Current Opinions in Neurobiology, 4*, 812–822.
- Kveraga, K., Ghuman, A., & Bar, M. (2007). Top-down predictions in the cognitive brain. *Brain and Cognition, 65*, 145–168.
- Lamme, V., & Roelfsema, P. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences, 23*, 571–579.

- Livingston, K., Andrews, J., & Harnad, S. (1998). Categorical perception effects induced by category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 732–753.
- Lupyan, G. (2007). Reuniting categories, language, and perception. In D. McNamara & J. Trafton (Eds.), *Proceedings of the Twenty-Ninth Annual Meeting of the Cognitive Science Society* (pp. 1247–1252). Austin, TX: Cognitive Science Society.
- Lupyan, G. (2008a). The conceptual grouping effect: Categories matter (and named categories matter more). *Cognition*, *108*, 566–577.
- Lupyan, G. (2008b). From chair to "chair:" A representational shift account of object labeling effects on memory. *Journal of Experimental Psychology: General*, *137*, 348–369.
- Mack, M.L., Gauthier, I., Sadr, J., & Palmeri, T.J. (2008). Object detection and basic-level categorization: Sometimes you know it is there before you know what it is. *Psychonomic Bulletin & Review*, *15*, 28–35.
- McCotter, M., Gosselin, F., Maccabee, P.J., & Schyns, P. (2005). The use of visual information in natural scenes. *Visual Cognition*, *12*, 938–953.
- McMurray, B., Aslin, R.N., Tanenhaus, M.K., Spivey, M., & Subik, D. (2008). Gradient sensitivity to within-category variation in words and syllables. *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 1609–1631.
- Mitterer, H., Horschig, J.M., Musseler, J., & Majid, A. (2009). The influence of memory on perception: It's not what things look like, it's what you call them. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 1557–1562.

- Norris, D., McQueen, J.M., & Cutler, A. (2000). Merging information in speech recognition: Feedback is never necessary [Target article and commentaries. *The Behavioral and Brain Sciences*, 23, 299–370.
- Notman, L., Sowden, P., & Ozgen, E. (2005). The nature of learned categorical perception effects: A psychophysical approach. *Cognition*, 95, B1–B14.
- Posner, M. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Erlbaum.
- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–365.
- Quinn, P., Eimas, P.D., & Tarr, M.J. (2001). Perceptual categorization of cat and dog silhouettes by 3- to 4-month-old infants. *Journal of Experimental Child Psychology*, 79, 78–94.
- Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience*, 3, 1199–1204.
- Schyns, P., & Oliva, A. (1999). Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition*, 69, 243–265.
- Schyns, P., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 681–696.
- Spivey, M. (2008). *The continuity of mind*. Oxford, England: Oxford University Press.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327–352.
- VanRullen, R., & Thorpe, S.J. (2001). The time course of visual processing: From early perception to decision-making. *Journal of Cognitive Neuroscience*, 13, 454–461.
- Winawer, J., Witthoft, N., Frank, M., Wu, L., Wade, A., & Boroditsky, L. (2007). Russian blues reveal effects of language on color discrimination. *Proceedings of the National Academy of Sciences, USA*, 104, 7780–7785.

Figure Captions.

Fig. 1.

Experimental design for Experiments 1 through 3. The two letters on each trial were the same, physically different but from the same conceptual category, or physically different and from different conceptual categories. These letters were presented either simultaneously or sequentially. On the sequential trials, the first stimulus remained visible during a variable delay (stimulus onset asynchrony, or SOA) before onset of the second stimulus.

Fig. 2.

Results from (a) Experiment 1 (upright letters) and (b) Experiment 2 (rotated letters). The graphs show mean reaction times on *different* trials as a function of stimulus onset asynchrony (SOA), separately for within-category and between-category trials. Error bars represent ± 1 *SE* of the mean difference score for a given SOA.

Fig. 3.

Mean reaction times on *different* trials in Experiment 3 as a function of stimulus onset asynchrony (SOA), separately for within-category and between-category trials. Error bars represent ± 1 *SE* of the mean difference score for a given SOA.

Fig. 4.

Probability density functions of participants in Experiments 1 and 3. Each point represents the normalized frequency of a peak category effect (reaction time for within-category trials minus reaction time for across-category trials) for a particular stimulus onset asynchrony (SOA).

Fig. 5.

The stimuli used in Experiment 4 and the mean time to classify each as a “cat” or a “dog.” C1 and D1 served as “typical” stimuli. C3 and D3 served as “atypical” stimuli.

Fig. 6.

Mean reaction times on *different* trials in Experiment 4 as a function of stimulus onset asynchrony (SOA), separately for within-category (cat_a-cat_b, dog_a-dog_b) and between-category (cat-dog) trials. Error bars represent ± 1 *SE* of the mean difference score for a given SOA.

Figure 1.

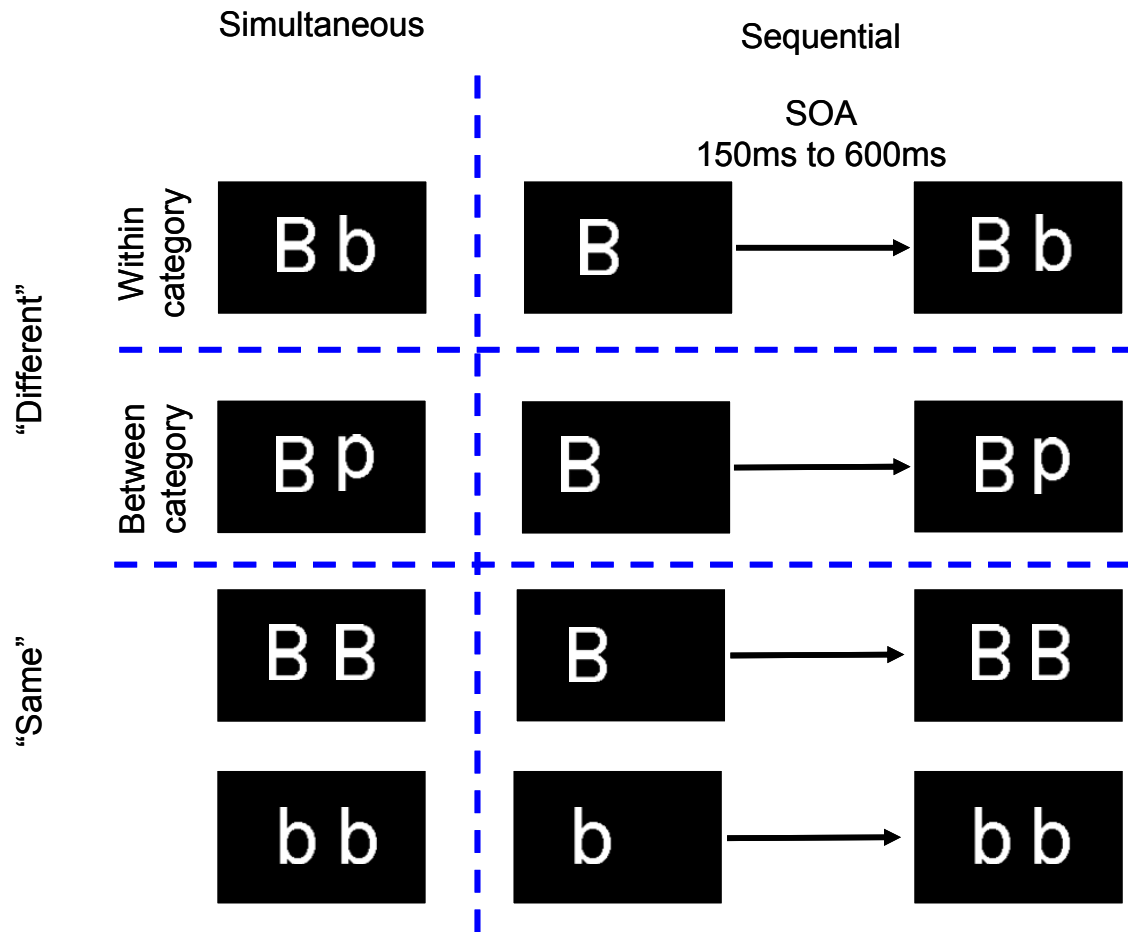



Figure 2.

a



b




Figure 3.




Figure 4.

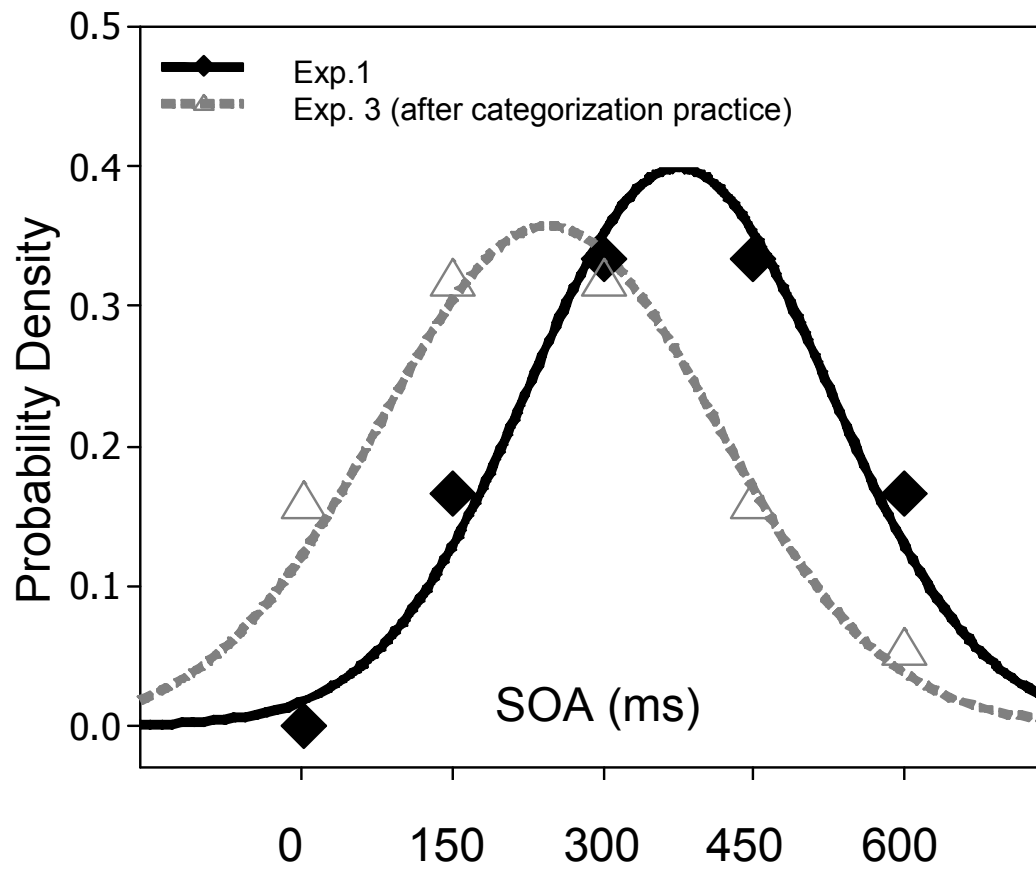


Figure 5.



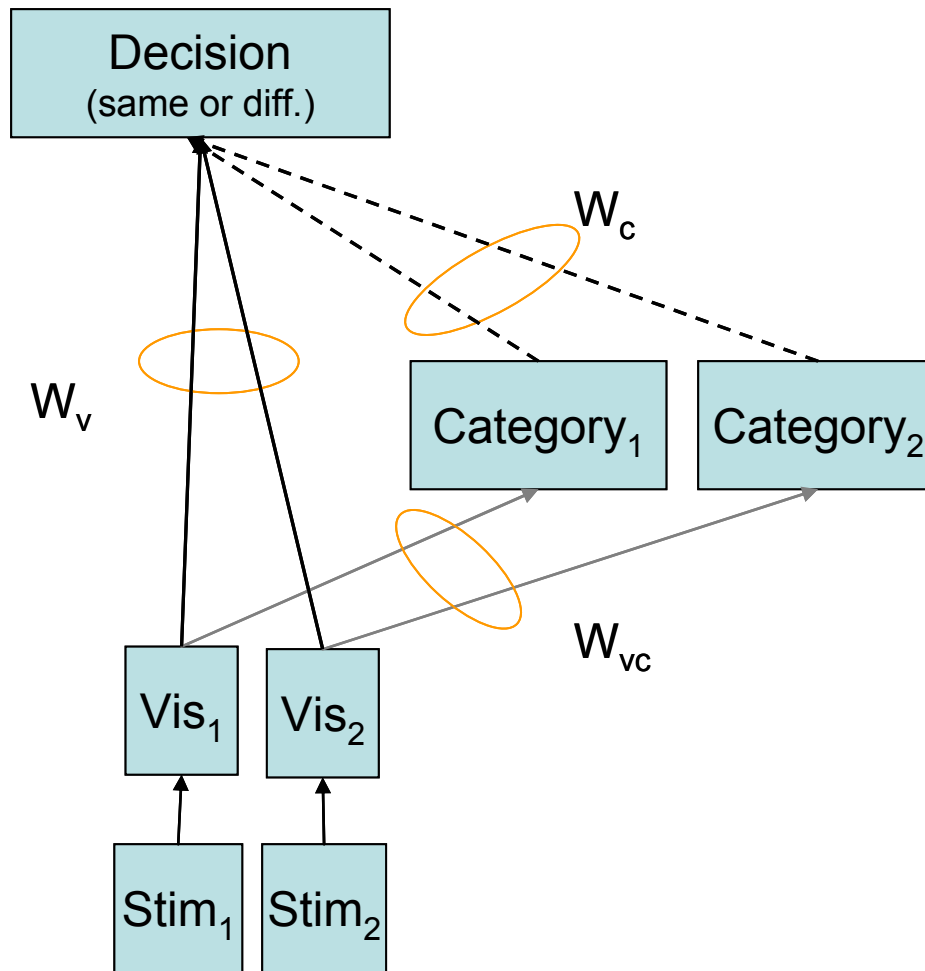


Figure 6.



A further explanation of why a decision-level model cannot account for a category effect that varies with the onset delay of the second stimulus.



Consider the model above (fashioned after the Merge model of Norris, McQueen, & Cutler, 2000; see also Mitterer, Horschig, Musseler, & Majid, 2009). On this model, the category effect – slower responses for within-category than between-category trials – is modeled in terms of competition between a response based on a comparison of the visual features (solid black lines)

and the response based on a comparison of the category membership (dashed lines). That is, in this model visual representations are *not* penetrated by conceptual factors. A greater category effect can be produced by increasing W_c , and/or decreasing W_v . Critically, neither the visual weights nor the category weights contribute meaningfully to the decision until *both* stimuli are present.

Now, let us run through this model for the simultaneous presentation case. Let $\text{Stim}_1 = B$ and $\text{Stim}_2 = b$. After a visual feature analysis, the stimuli activate their appropriate category representations, which are identical for these two stimuli. A decision based on the visual features would produce a “different” response, while a decision based on the conceptual features would produce a “same” response. The degree to which the final decision is affected by the (misleading) information from the category layer depends on the relative contributions of W_v and W_c . In Experiment 1, no reliable effect was found for simultaneous presentations, so let us set W_v and W_c to values that yield a minimal contribution of the category layer to the final decision. The question now is whether there is any way to increase the influence of the category layer on the decision for sequential judgments.

Let us run through a sequential trial. Stim_1 is presented and automatically categorized. Sometime during or after the categorization process (depending on the length of the SOA and strength of W_{vc}), Stim_2 is presented and starts being categorized. Critically, the SOA is irrelevant for the category effect because neither the visual layers nor the category layers can inform the decision until *both* letters are present. The categorization process can be slowed down/inhibited by decreasing W_{vc} . The category effect will be appropriately decreased: the decision will now be

reached based on the visual features because the category information has not had a chance to propagate. But notice that this decrease would affect simultaneous and sequential trials equally because in this model, a single item (categorized or not) does not affect the decision process.

A purely feed-forward model can be made to work by adding an additional layer (call it V-C) which combines inputs from both the visual and category information. V-C would in turn feed into the decision layer. Note, however, that this is no longer a decision-level model because category-level representations are altering visual representations rather than decision-level representations

References:

- Mitterer, H., Horschig, J. M., Musseler, J., & Majid, A. (2009). The Influence of Memory on Perception: It's Not What Things Look Like, It's What You Call Them. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(6), 1557-1562.
- Norris, D., McQueen, J. M., & Cutler, A. (2000). Merging information in speech recognition: feedback is never necessary. *The Behavioral and Brain Sciences*, 23(3), 299-325; discussion 325-370.