



ELSEVIER

Contents lists available at ScienceDirect

## Journal of Memory and Language

journal homepage: [www.elsevier.com/locate/jml](http://www.elsevier.com/locate/jml)

## Visual interference disrupts visual knowledge



Pierce Edmiston\*, Gary Lupyan

Department of Psychology, University of Wisconsin–Madison, United States

## ARTICLE INFO

## Article history:

Received 17 April 2015

revision received 2 July 2016

## Keywords:

Concepts

Representations

Visual knowledge

Word-recognition

Visual long-term memory

## ABSTRACT

We show that visual interference impairs people's ability to make use of visual knowledge. These results provide strong evidence that making use of stored visual knowledge—long-term memory of what things look like—depends on perceptual mechanisms. In the first set of studies, we show that presenting visual noise patterns during or after hearing a verbal cue greatly reduces the effectiveness of the cue on a simple visual discrimination task. In the second experiment, participants were tasked with answering questions about visual features of familiar objects, e.g., verifying that tables have flat surfaces. Accuracy in answering visual, but not encyclopedic questions was reduced when viewing colorful noise patterns. This result is most parsimoniously explained by positing that judgments required activation of visual representations that were being interfered with when viewing irrelevant patterns. Although much of our conceptual knowledge may abstract away from perceptual details, knowledge of what things look like appears to be represented in a visual format.

© 2016 Published by Elsevier Inc.

## Introduction

Much of what people know about objects comes from visual experience (Cree & McRae, 2003; Hoffman & Lambon Ralph, 2013). For example, our knowledge that alligators are animals with sharp teeth and long tails comes at least in part from seeing alligators, even if only in pictures. Yet, many have argued that although we learn many things *via* our sense of vision, the visual knowledge itself (also called visual long-term memory) is amodal—represented and accessed independently of perceptual processes (Caramazza, Hillis, Rapp, & Romani, 1990; Pinker, 1994; Pylyshyn, 1986; Tulving, 1972). Apparently contradicting this amodal view are studies showing that asking people to think about what things look like elicits patterns of neural activity that overlap with patterns evoked by

actual visual processing (Amsel, Urbach, & Kutas, 2014; Martin, 2007; Simmons et al., 2007). At present, however, there is little direct evidence for the *causal* involvement of visual processes in making use of visual knowledge (Mahon & Caramazza, 2008; Papeo, Pascual-Leone, & Caramazza, 2013). It remains possible that the perceptual activation measured in such tasks is epiphenomenal or reflects explicit mental imagery (Albers, Kok, Toni, Dijkerman, & de Lange, 2013; Kosslyn, Ganis, & Thompson, 2001; Naselaris, Olman, Stansbury, Ugurbil, & Gallant, 2015). If, however, it can be shown that disrupting visual processes disrupts visual knowledge (and *specifically* visual knowledge),<sup>1</sup> a parsimonious conclusion would be that this visual knowledge was at least partly constituted by the now disrupted representations.

To establish whether visual knowledge and visual perception rely on common mechanisms, we interfered with

\* Corresponding author. Address: Department of Psychology, University of Wisconsin–Madison, 1202 W. Johnson St., Madison, WI 53706, United States.

E-mail addresses: [pedmiston@wisc.edu](mailto:pedmiston@wisc.edu) (P. Edmiston), [lupyan@wisc.edu](mailto:lupyan@wisc.edu) (G. Lupyan).

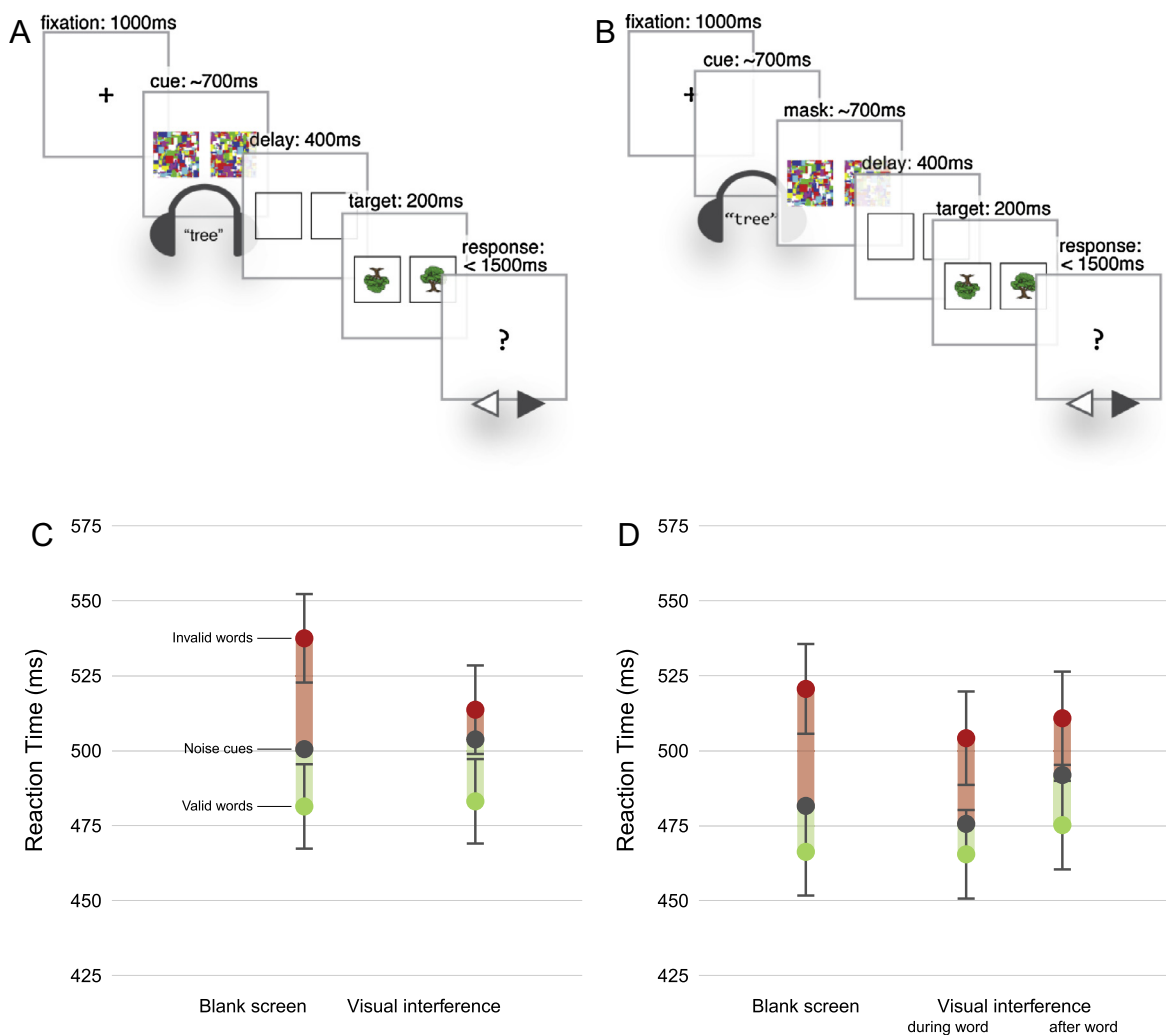
<sup>1</sup> Our claims do not require visual interference to completely abolish the ability to make use of visual knowledge. The key requirement is that visual interference should impair visual—and specifically visual—knowledge.

visual perception and measured the impact of this interference on simple behavioral tests of visual knowledge. To foreshadow the results: In Experiments 1A–1C, visual interference reduced a robust verbal cueing effect where hearing the name of an object speeds its visual recognition. These results show that the effect of hearing verbal cues on visual recognition (e.g. the benefit of hearing “dog” on recognizing a picture of a dog) results from the label activating visual representations that can be disrupted by irrelevant visual input. In Experiment 2, we extend these results to a completely linguistic domain and find that visual interference impaired accuracy in verifying statements concerning visual properties of common categories, e.g., that tables have flat surfaces, but left encyclopedic and otherwise nonvisual knowledge, e.g., that tables are furniture, unaffected. These results provide compelling

evidence that perceptual representations are at least partly constitutive of our knowledge of what things look like.

### Experiment 1A: Visual interference disrupts verbal cueing of visual recognition

Hearing a name affects visual recognition of the named category. For example, in the task depicted in Fig. 1A, hearing “tree” prior to seeing pictures of an upright and an upside-down tree makes it easier to recognize which tree is upright and which is upside-down (a validity advantage) while slowing performance on judging the orientation of another object, e.g., a car (an invalidity cost) relative to baseline trials on which no word is heard (Lupyan & Thompson-Schill, 2012).



**Fig. 1.** Visual interference reduces verbal cueing of visual knowledge. (A) A sample trial from Experiment 1A in which visual interference is presented during a valid verbal cue prior to deciding which of the two images was upright. (C) Results of Experiment 1A when interference was presented during the auditory word cues. Valid cues improved performance (size of green bars) and invalid cues (e.g., “dog” preceding tree pictures) impaired performance (size of red bars) relative to baseline in which the verbal cue was replaced with uninformative white noise. The total cueing effect (size of full bar) was reduced by visual interference. (B) A sample trial from Experiment 1B which included trials in which the visual interference was delayed until the offset of the auditory cue. (D) Results from Experiment 1B when interference was presented during or after the word cues. The total cueing effect was reduced to a similar extent even when the interference was delayed. Error bars show  $\pm 1$  SE of coefficient estimates.

This type of cueing effect can be explained in two ways. If the knowledge of what objects look like is represented outside the visual system in an amodal format, then the verbal cueing effect may affect performance by activating abstract, non-perceptual representations such as lists of category features or nodes in a semantic network to which the target picture is then matched. Precisely such a mechanism is thought to underlie classic semantic priming effects (e.g., being faster to verify that “doctor” is a word after hearing the semantically related “nurse”) (Caramazza et al., 1990; Collins & Loftus, 1975; Firestone & Scholl, 2015). On an alternative view (Barsalou, Simmons, Barbey, & Wilson, 2003; Lupyan & Thompson-Schill, 2012), the word activates a perceptual representation of the object, aiding in processing subsequently presented visual inputs: in this case, helping to discriminate the target in its canonical orientation from a target in a non-canonical orientation. On this latter view, the visual knowledge of e.g., what dogs look like is being instantiated by the word “dog” *within* the perceptual system. While both accounts predict the same pattern of behavioral performance, recent electrophysiological evidence from ERPs (Boutonnet & Lupyan, 2015) has found that verbal cues modulate the early visual response component (P100) in a way that correlated with visual recognition performance following the verbal cue. This finding lends support to the idea that labels activate perceptual information (see also Lupyan & Ward, 2013; Meteyard, Zokaei, Bahrami, & Vigliocco, 2008).

The use of visual interference, however, makes for an especially strong test of the hypothesis that the changes in visual recognition after processing a verbal label are due to the label activating perceptual states. We hypothesized that if the cueing effect results from the activation of visual representations then visual interference occurring prior to the onset of the target images would reduce or eliminate the effect of the verbal cue on subsequent image recognition.

## Methods

### Participants

We recruited 25 University of Wisconsin-Madison undergraduates for course credit.

### Materials

We selected 38 color line-drawings (Rossion & Pourtois, 2004) on the basis of their high imageability scores. Selected drawings pictured familiar animals and objects (e.g., alligator, bottle, giraffe, piano, snowman, tree, zebra). Images and accompanying masks subtended  $\sim 8^\circ \times 8^\circ$  visual angle. Verbal cues (the names of the objects: “alligator”, “bottle,” etc.) were recorded by a female speaker at a normal reading rate, and ranged from 500 ms to 1000 ms. Each participant was randomly assigned to receive 10 of the 38 possible categories.

### Procedure

On each trial, participants heard an auditory cue followed by the presentation of two pictures varying only in their orientation and had to determine which picture was correctly oriented (Fig. 1A). All trials began with a central

fixation cross displayed for 1000 ms, after which an auditory cue was played through the headphones. On 50% of all trials, the cue was a word. On the remaining trials, which provided a measure of baseline performance, participants heard uninformative white noise (700 ms duration to match the average duration of the verbal cues).

Of the word-cued trials, the word validly identified the subsequently presented pictures 75% of the time (e.g., “zebra” followed by two zebras). On the remaining 25% of the word-cued trials, the cue was invalid, e.g., participants might hear “dog” followed by two alligators.

On a random 50% of all trials, participants saw visual interference in the form of colorful rectangles with sizes, colors, and positions alternating at  $\sim 60$  Hz. Visual interference was presented during the auditory cue. After auditory cue offset, the visual interference patterns were replaced by two rectangular frames that marked the location of the to-be-presented images. After approximately 400 ms, the two rectangles were filled with the target images. One of the images was upright and the other upside down. The images were visible for 200 ms after which they were replaced by a central question mark prompt indicating that the participant should press the key corresponding to the side (left or right) that displayed the upright image. Participants responded using the keyboard. If a participant did not respond in 1500 ms, a timeout prompt was displayed, and the trial was discarded from analyses.

Participants began by completing 10 practice trials. During these practice trials only, participants received auditory feedback (a bleep sound for correct answers and a buzz for errors). Participants completed 320 test trials, with three evenly spaced breaks. Trials were pseudo-randomized to minimize the number of times the same picture target appeared on back-to-back trials.

### Exclusion criteria

One of the participants did not complete the experiment and the data were excluded. Data from three additional participants were excluded due to slow response times and exceedingly high error rates ( $\sim 15\%$ ).

### Statistical modeling

Response times (ms) for correct responses on the test trials were fit with mixed linear regression models (lme4 package; Bates, Maechler, Bolker, & Walker, 2014) in R (R Core Team, 2014). We report coefficients corresponding to the cueing effect—the difference in RTs between valid, noise, and invalid cues—with corresponding 95% confidence intervals. In the models, we allowed the overall response time to vary by participant (random intercepts). The verbal cueing effect—the difference between response times on valid and invalid trials relative to the noise trials—was estimated in the regression models by creating two orthogonal contrasts: one which contrasted the effect of a valid cue versus an invalid cue, and another which contrasted the effect of verbal cues versus nonverbal (noise) cues. Both contrasts were allowed to interact with the visual interference condition. Degrees of freedom for multi-level regression models (and subsequent significance tests) were estimated using the lmerTest package (Kuznetsova, Brockhoff, & Christensen, 2014).

## Results and discussion

As expected, the speed at which participants discriminated upright objects was improved by hearing the name of the object prior to seeing it, as indicated by a 19 ms reduction in response times relative to baseline trials in which a noise cue was heard instead,  $-19$  ms, 95% CI  $[-27, -11]$ ,  $z = -4.6$ ,  $p < 0.001$ . Conversely, if the objects that were presented were different than the named object, performance was impaired by 37 ms relative to noise cue baseline trials, 95% CI  $[24, 50]$ ,  $z = 5.8$ ,  $p < 0.001$ . The combination of valid cue facilitation and invalid cue impairment resulted in a total cueing effect of 56 ms, 95% CI  $[44, 69]$ ,  $z = 8.8$ ,  $p < 0.001$  on trials in which no visual interference was presented (Fig. 1C). Cues did not affect response accuracy,  $-0.59$  log-odds, 95% CI  $[-1.39, 0.20]$ ,  $z = -1.5$ ,  $p = 0.14$ , likely due to the already very high accuracy level: the average participant made fewer than 11 errors out of 320 trials ( $M = 3.4\%$ ).

On trials with visual interference, the cueing effect in RTs was reduced by 25 ms, a 45% reduction, 95% CI  $[-43, -8]$ ,  $z = 2.9$ ,  $p = 0.004$  (Fig. 1C). The effect of visual interference was greater on verbal cue trials than it was on noise cue trials, a 14 ms difference, 95% CI  $[2.8, 26]$ ,  $z = 2.4$ ,  $p = 0.015$ , indicating that visual interference by itself did not disrupt baseline image recognition. Looking only at the noise cue trials, visual interference was not found to slow baseline performance, a 3 ms difference, 95% CI  $[-4, 11]$ ,  $z = 0.83$ ,  $p = 0.41$ . Post-hoc analyses revealed that the reduction in the cueing effect was driven primarily by a 24 ms improvement in response speed on invalid cue trials, 95% CI  $[8, 40]$ ,  $z = 2.9$ ,  $p = 0.004$ . That visual noise presented during invalid-cue trials improved absolute performance as expected if the noise disrupts incongruent visual content activated by hearing the word, content that otherwise interferes with locating the upright object.

Baseline error rates were not increased by visual interference on noise cue trials ( $M_{baseline} = 3.4\%$ ,  $M_{interference} = 3.2\%$ ),  $-0.07$  log-odds, 95% CI  $[-0.42, 0.29]$ ,  $z = -0.37$ ,  $p = 0.71$ . Visual interference did not affect error rates on the invalid cue trials, ( $M_{baseline} = 4.0\%$ ,  $M_{interference} = 3.6\%$ ),  $-0.11$  log-odds, 95% CI  $[-0.80, 0.56]$ ,  $z = -0.33$ ,  $p = 0.74$ , but visual interference did increase errors 1.6-fold on valid cue trials ( $M_{baseline} = 2.6\%$ ,  $M_{interference} = 4.2\%$ ), 0.48 log-odds, 95% CI  $[0.07, 0.90]$ ,  $z = 2.3$ ,  $p = 0.02$ .

The results show that viewing colorful patterns—a form of visual interference—disrupted a verbal cueing effect without affecting baseline task performance. The most parsimonious explanation is that the visual knowledge activated by hearing a word is in part represented as visual representations that can be partly disrupted by simply viewing noise patterns.

### Experiment 1B: Delayed visual interference still disrupts verbal cueing of visual recognition

An alternative explanation of the results obtained in Experiment 1A is that visual interference simply distracted people from hearing (or otherwise processing) the auditory word cues. We sought to rule out this possibility through two additional experiments (1B, 1C). In 1B we compared

the effect of visual interference when it occurred *during* the presentation of the word cues to when it occurred *after* the presentation of the auditory word cues. If the colorful patterns simply distracted people from hearing the word, we expected the interference to be considerably less effective when presented after the word. If, however, visual interference disrupts the activation of visual knowledge, it should have a disruptive effect even if presented shortly after the word because it is predicted to disrupt an activation trajectory unfolding on a timescale of  $\sim 1.5$  s (Lupyan & Spivey, 2010; Lupyan & Thompson-Schill, 2012).

## Methods

### Participants

We recruited 38 University of Wisconsin-Madison undergraduates who received course credit for participation.

### Materials and procedure

Experiment 1B was identical to 1A except the visual interference trials were divided into concurrent trials for which visual interference was present *during* the auditory cue, as in Exp. 1A, and sequential trials for which visual interference was present *after* the offset of the word cues.

To allow for sequential visual interference while keeping the time between auditory cue offset and picture onset the same for all trials, the time between all auditory cues and the picture targets was increased from 400 ms in Experiment 1A to  $\sim 1100$  ms in Experiment 1B.

## Results and discussion

There was again a large cueing effect in RTs observed on the trials without any visual interference, 54 ms, 95% CI  $[45, 64]$ ,  $z = 11$ ,  $p < 0.001$ . As in Experiment 1A, no cueing effect in response accuracy was observed,  $-0.14$  log-odds, 95% CI  $[-0.66, 0.34]$ ,  $z = -0.57$ ,  $p = 0.57$ .

Visual interference during or after the verbal cue reduced the cueing effect in RTs by 32% (17 ms), 95% CI  $[3, 31]$ ,  $z = -2.5$ ,  $p = 0.01$ . Delaying the onset of the visual interference until after hearing the cue did not lessen the impact of visual interference on the cueing effect,  $M_{during} = 16$  ms,  $M_{after} = 19$  ms, a 3 ms difference, 95% CI  $[-16, 22]$ ,  $z = 0.31$ ,  $p = 0.75$ . That is, visual interference reliably reduced the cueing effect (by 19 ms) even when its presentation was delayed until after the offset of the verbal cue, 95% CI  $[2, 35]$ ,  $z = -2.2$ ,  $p = 0.028$  (Fig. 1D). This shows that visual interference did not distract participants from listening to the word cues. As in Experiment 1A, the reduction in the cueing effect was driven primarily by an improvement in response speed on invalid cue trials,  $M_{no\ interference} = 520$  ms,  $M_{interference} = 507$  ms, a 13 ms difference, 95% CI  $[0, 26]$ ,  $z = 1.9$ ,  $p = 0.054$ .

Visual interference did not affect error rates on noise cue trials,  $-0.10$  log-odds, 95% CI  $[-0.39, 0.18]$ ,  $z = -0.86$ ,  $p = 0.46$ , on valid cue trials,  $-0.30$  log-odds, 95% CI  $[-0.66, 0.06]$ ,  $z = -1.7$ ,  $p = 0.10$ , or on invalid cue trials, 0.24 log-odds, 95% CI  $[-0.34, 0.84]$ ,  $z = 0.82$ ,  $p = 0.41$ . These results did not differ by interference type (during or after),  $p$ 's  $> 0.32$ , except for on noise cue trials where errors were

more likely on delayed interference trials, 0.41 log-odds, 95% CI [0.00, 0.84],  $z = 2.0$ ,  $p = 0.047$ .

The results of experiment 1B show that visual interference does not disrupt the verbal cueing effect by impairing the initial encoding of a verbally presented cue. If it did, visual interference should have been considerably less effective when presented after the cue, but this was not the case. These results further support our hypothesis that visual interference disrupts the verbal cueing effect by affecting visual representations. A remaining possibility is that visual interference disrupts short-term memory for the verbal cue, independent of whether the meaning of the cue is represented in a visual format. We tested this possibility in Experiment 1C.

### Experiment 1C: Visual interference does not disrupt memory for the verbal cue

We conducted Experiment 1C to rule out the possibility that visual interference disrupts performance on any task requiring the processing of verbal cues by, for example, disrupting short-term memory of what the just-presented verbal cue was. If this were the case, then the reduction in the cueing effect observed in Experiments 1A–1B may again be due to the effect of visual interference on processing of the verbal cue, rather than due to visual interference affecting visual representations constitutive of visual knowledge.

In Experiment 1C, participants performed the cued-orientation judgment task with/without visual interference described in Experiment 1A, but on 25% of the trials they were instead shown a printed word and had to respond whether it was the same as the word they had heard earlier in the trial. We expected visual interference to again disrupt the verbal cueing effect, but not to have an effect on word repetition detection performance, as detecting the repetition of a verbal label does not require the visual representations hypothesized to be the cause of the verbal cueing effect.

### Methods

#### Participants

28 University of Wisconsin-Madison undergraduates were recruited for participation in Experiment 1C.

#### Materials and procedure

Experiment 1C was identical to Experiment 1A except that on a subset of trials (25%) instead of two pictures, a single word appeared in one of the two target boxes (left/right of fixation). Participants pressed the *up arrow* key if the printed word was the same as the word they heard previously in the trial or the *down arrow* if the printed word was different from the cue word they just heard. Of these word repetition trials, 75% of the time the printed word was a repetition of the auditory cue.

Participants were initially allowed to respond with any of the four arrow keys on each trial but often made response key errors (e.g., responding left/right on a word repetition trial). After initial data collection ( $N = 14$ ), we

prevented left or right responses on word repetition trials from registering as valid responses.

In Experiment 1C, participants heard a word on every trial, i.e., we excluded the baseline, noise trials presented in Experiments 1A–1B. This allowed participants to complete the experiment in approximately the same amount of time as Experiment 1A, and prevented participants from possibly inferring after hearing a noise cue that the trial was an orientation discrimination trial.

### Results and discussion

In a replication of Experiments 1A and 1B, we observed a reliable cueing effect on the trials without any visual interference, 71 ms, 95% CI [57, 86],  $z = -9.7$ ,  $p < 0.001$ , and visual interference reduced this cueing effect by 39% (a 28 ms reduction), 95% CI [9, 48],  $z = 2.8$ ,  $p = 0.005$  (see Fig. 2, left panel).

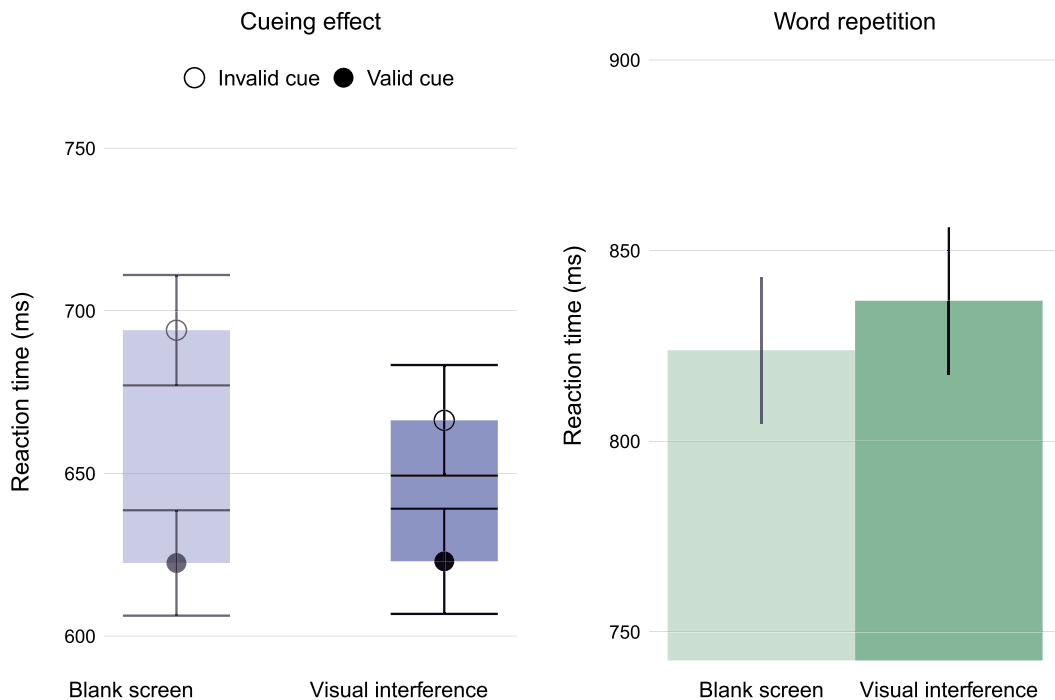
In contrast, visual interference did not have a detrimental effect on the speed, 13 ms, 95% CI [-7, 32],  $z = 1.3$ ,  $p = 0.19$ , or accuracy, 0.14 log-odds, 95% CI [-0.20, 0.49],  $z = 0.82$ ,  $p = 0.41$ , with which repeated words were detected.

To further rule out the possibility that the effect of visual interference on the verbal cueing effect was related to the effect of visual interference on the word repetition task, we calculated the effect of visual interference for each participant on the orientation-discrimination trials and word-repetition trials, and evaluated the relationship between visual interference effects across participants. Across participants the effect of visual interference on the orientation-discrimination trials was not predictive of the effect of interference effect on word-repetition repetition trials,  $b = 0.035$ , 95% CI [-0.35, 0.42],  $p = 0.85$ , providing further evidence that the effect of visual interference on the effect of the verbal cue was not caused by people forgetting or not attending to the cue.

### Discussion

In Experiments 1A–1C, the verbal cueing of visual recognition was significantly reduced by visual interference. We interpret this result as showing that the verbal cueing of visual recognition depends on the activation of visual representations which are interfered with by the flashing rectangle patterns.

Together, Experiments 1A–1C rule out the main alternative explanations of the disruptive effect of visual interference on verbal cueing. The possibility that visual interference might be generally distracting or disrupting was ruled out by a significant interaction between the effect of verbal interference as a function of cue-type: visual interference only affected performance on trials containing verbal cues, and not on noise cue trials. The possibility that visual interference reduced the cueing effect by simply distracting participants from listening to the verbal cues was ruled out by Experiment 1B which showed that the effect of verbal interference was nearly identical whether it occurred during the word or immediately after the word. The remaining possibility that visual interference might have made participants forget the verbal cue



**Fig. 2.** Results of Experiment 1C designed to rule out the possibility that visual interference disrupts memory for the verbal cue. Visual interference reduced the cueing effect (left). The same visual interference did not disrupt memory for the name of the object as measured by reaction times (right) or accuracies (not shown; see text), presumably because detecting a repeated word does not require visual representations disrupted by the visual interference. Error lines are  $\pm 1$  standard error of the model estimates (see text for model descriptions).

was ruled out by Experiment 1C which showed the effect of verbal interference to be selective to disrupting the verbal cueing of visual recognition as it did not disrupt one's ability to detect the cued word presented visually.

The finding that *visual* interference disrupts *verbal* cueing of visual knowledge is most parsimoniously explained by positing that the visual knowledge activated by verbal labels is represented (at least in part) in a visual format. The alternative view—that visual knowledge is amodal and distinct from visual perception—does not provide a mechanism by which viewing task-irrelevant visual interference can selectively interfere with the activation of visual knowledge by verbal cues.

Experiments 1A–1C show that visual interference specifically disrupted the use of verbal cues on a visual task—a conceptually mediated effect. Yet, the task of identifying upright images is still a visual task. The hypothesis we are advocating—that visual knowledge is, at least partially, represented in a perceptual format—does not require the task itself to be visual. If visual interference disrupts visual knowledge, then it may disrupt people's ability to answer even simple yes or no questions about what common objects and animals look like. This stronger test of our hypotheses was investigated in Experiment 2.

### Experiment 2: Visual interference disrupts verification of visual properties

We measured the effect of visual interference on people's ability to answer yes or no questions about the

properties of familiar objects. Our prediction was that visual interference would impair visual knowledge, as indicated by lower accuracy and longer latency when answering questions about visual properties. Such disruptions cannot be easily accommodated by an amodal view of visual knowledge because on this view assessing the truth value of verbal propositions does not require the activation of any perceptual representations.

### Methods

#### Participants

A total of 149 University of Wisconsin-Madison undergraduates participated in exchange for course credit. The original sample had 25 participants. We next conducted a replication of the experiment with 79 additional participants. After analyzing the results of the original experiment and the replication attempt, we recruited 45 additional participants for a preregistered replication with minor changes to the methods, described below. Our main hypotheses do not differ across these experiment runs, and the results reported below include all recruited participants. The exercise of conducting the same experiment in the same lab across multiple years has proven painfully educational. An analysis of the experiment including differences between runs and the motivations for the final, preregistered design is provided in the Supplemental Materials.

#### Materials

We selected 30 target objects from [McRae, Cree, Seidenberg, and McNorgan \(2005\)](#) with the criterion that

each object had 4–5 listed visual features (visual form, visual color) and nonvisual features (e.g., functional, tactile, encyclopedic) equated on production frequency. The names of the objects (“alligator”, “bottle”, etc.) were recorded by a female speaker at a normal reading rate, and ranged in duration from 500 ms to 1000 ms. We then constructed 538 feature–object pairings (propositions) to use as questions in the experiment, e.g., Does it have a hump? → camel. For half, the correct answer was “yes.” For the remaining, the correct answer was “no”, e.g., Does it have seeds? → camel.

### Norming

The key manipulation in this study—the “visualness” of a question—was derived from the coding of feature–associations norms by [McRae et al. \(2005\)](#) in which features are coded by the authors as being visual, auditory, encyclopedic, etc. Of course some visual features are more visual than others. For example although having handlebars and being made of metal are both coded in [McRae et al.’s](#) norms as *visual* features of bicycles, having handlebars is arguably a more visually salient property than being made of metal. In designing the experiment the propositions are grouped into visual and nonvisual questions, and we collected norming data to verify that the visual questions indeed required more visual knowledge than the nonvisual questions as rated by independent samples of people. Participants in the norming study also rated the propositions for a number of plausible confounds, including the amount of encyclopedic knowledge and the degree to which verifying the proposition required abstract knowledge, as well as assessing the difficulty and the truth-value of each proposition.

Ratings were provided by participants naïve to the visual interference manipulation recruited via Amazon’s Mechanical Turk. Each participant was presented with 20 object–feature pairs (propositions), e.g., dog: Does it have four legs? 145 participants provided the ratings used to estimate the *amount of visual knowledge* as well as the *amount of encyclopedic knowledge*, the *truth-value* and the *difficulty* of each proposition used in Experiment 2. We subsequently recruited an additional 120 participants to collect ratings of the *abstractness* of each proposition.

*Amount of visual knowledge* was assessed by asking participants “How much is it necessary to picture the object in your mind’s eye?” in order to verify whether the provided statement was true. Participants responded using a 5-point Likert scale anchored at (1) *None* and (5) *As much as possible*. As expected, visual questions were rated as requiring more visual knowledge than nonvisual questions,  $M_{\text{visual}} = 2.58$ ,  $M_{\text{nonvisual}} = 2.28$ ,  $p < 0.001$ .

*Amount of encyclopedic knowledge* was assessed by asking a parallel question, “How much did you rely on facts that you might read in an encyclopedia or school textbook?” to answer the question. Participants responded using the same 5-point Likert scale with anchors at (1) *None* and (5) *As much as possible*. Nonvisual questions were rated as requiring more encyclopedic knowledge than visual questions,  $M_{\text{visual}} = 1.91$ ,  $M_{\text{nonvisual}} = 2.13$ ,  $p < 0.001$ .

Nonvisual questions included both questions pertaining to features in other perceptual modalities (e.g., auditory

features) as well as more abstract domains (e.g., taxonomic distinctions). To measure whether the visual interference was sensitive to the dimension of *abstractness* independent of perceptual modality, we also had participants rate each proposition according to how likely it was that the information was learned through one’s senses, independent of whether those senses were visual, auditory, somatosensory, etc. Specifically, they were asked, “Is this something you can learn through your senses?” and responded on a scale from (1) *Definitely true* to (5) *Definitely not true*. Features that were not learned directly through the senses were considered more abstract. Overall, nonvisual questions were more abstract than visual questions,  $M_{\text{visual}} = 1.9$ ,  $M_{\text{nonvisual}} = 2.4$ ,  $p < 0.001$ , but abstractness was independent of *amount of visual imagery*,  $r = -0.043$ ,  $p = 0.32$ .

*Difficulty* was assessed by asking participants to rate the proposition for how difficult it was, subjectively, to verify, from (1) *Very Easy* to (5) *Very Difficult*. Visual and nonvisual questions did not differ in terms of rated difficulty,  $M_{\text{visual}} = 1.58$ ,  $M_{\text{nonvisual}} = 1.63$ ,  $p = 0.15$ . Controlling for question difficulty allows us to ensure that visual interference was selectively disrupting the activation of visual knowledge and not simply interfering with questions that were more difficult.

*Truth-value* was assessed by asking participants to answer the proposition (e.g., “lamp: Does it have a cord?”) on a 5-point Likert scale from (1) *Definitely No* to (5) *Definitely Yes*. Visual and nonvisual questions did not differ in overall truth value,  $M_{\text{visual}} = 2.9$ ,  $M_{\text{nonvisual}} = 3.0$ ,  $p = 0.42$ .

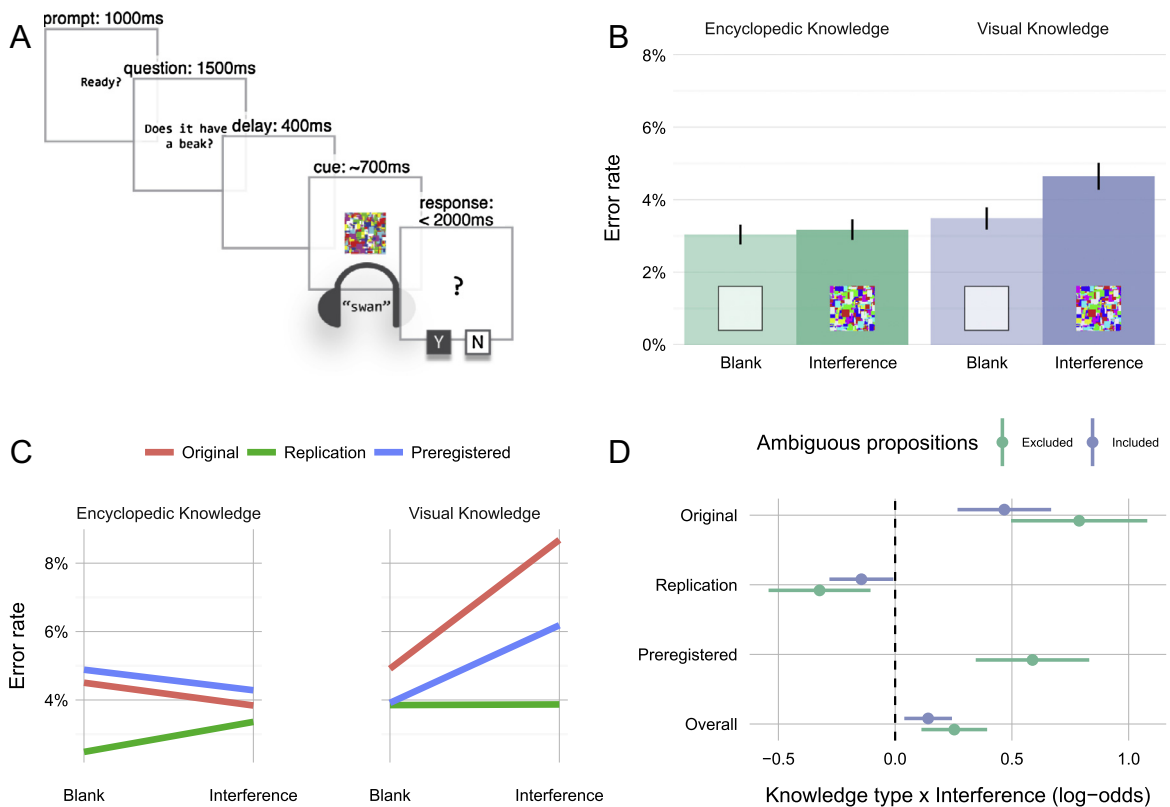
### Procedure

A sample trial with visual interference is shown in [Fig. 3A](#). Participants answered simple yes/no questions about familiar objects (e.g., whether swans have long necks) while looking at a blank screen (no interference) or while looking at colorful patterns (interference).

Trials began with 1000 ms “Ready?” prompt. The prompt was then replaced with the question (e.g., “Does it have a long neck?”) appearing at the center of the screen. After 1500 ms the question disappeared and after a 400 ms delay an auditory verbal cue (e.g., “swan”) was played. On a random 50% of trials, the auditory cue was accompanied by colorful rectangles (“Mondrians”) flashing at  $\sim 100$  Hz presented centrally subtending  $\sim 12^\circ \times 12^\circ$  of visual angle. We will refer to these trials as the visual interference trials.

Participants answered an equal number of visual and nonvisual questions based on the [McRae et al.’s \(2005\)](#) norms. Questions were randomly selected for each participant such that for 75% of the trials, the correct response was “yes”. After the offset of the final stimulus, participants saw a “?” prompt, and responded “yes” or “no” using the keyboard. If the participant did not respond within 2000 ms, a timeout prompt was displayed and the trial was excluded from analysis.

Participants began by completing 18 practice trials. During these practice trials only participants received auditory feedback (a bleep sound for correct answers and a buzz for errors). Participants completed approximately 200 test trials, with three evenly spaced breaks. Trials were pseudo-randomized such that participants were not asked questions



**Fig. 3.** Visual interference impairs visual and only visual knowledge. (A) A sample trial in which a yes/no question is asked in the presence of visual interference. Question type and interference condition were varied within-subjects. (B) Error rates by question type and interference condition. Questions were created from object feature norms (McRae et al., 2005). Nonvisual questions (functional, tactile, etc.) were selected to be equal in difficulty to questions about visual form and color. (C) Error rates across replication attempts. After running the original study and failing to replicate it, an improved experimental design was preregistered. (D) Knowledge type by interference interaction with and without ambiguous propositions. The preregistered experiment was conducted with ambiguous propositions removed. Error bars/bands show  $\pm 1$  SE of model estimates factoring out between-subject variance.

about the same category on back-to-back trials. Each participant saw a given question  $\rightarrow$  cue pairing only once.

#### Pilot study

We conducted a pilot study using the same materials described above, but with a slight variation in the procedure. In this pilot, participants heard the object cue prior to reading the question, e.g., “camel  $\rightarrow$  Does it have a hump?” The results of the pilot study, which are reported in the Supplemental Materials, motivated the design of Experiment 2 in which the feature preceded the object cue.

#### Preregistered replication

After conducting the original study and a failed attempt to replicate, a second replication effort was preregistered and administered <osf.io/9dy87>. We made two changes to the methods for the preregistered replication attempt, neither of which affected our primary hypothesis that visual interference would disrupt visual and only visual knowledge. First, question validity was changed from 75% to 50% in the preregistered replication. This was intended to increase the attentiveness of the participants by preventing them from having a ‘yes’ bias when responding. The second change related to the response window. In the original design, participants were prevented from responding until

after the offset of the final auditory cue, and as a result, reaction times were relatively fast ( $M = 335$  ms), suggesting that participants were waiting to provide a response. In fact, in the initial results visual interference was found to actually improve response times across all trial types,  $-13$  ms, 95% CI  $[-22, -5.1]$ ,  $p = 0.0016$ , suggesting to us that the offset of the visual interference might have unintentionally provided the participants a visual cue for the onset of the response window. To avoid this potential confound, in the preregistered replication participants were allowed to respond as soon as the auditory cue began playing.

#### Exclusion criteria

Of the 149 participants recruited for the experiment, 12 were excluded because of experiment error in the refresh rate of the monitor which altered the presentation rate of the visual interference. Two additional participants were excluded due to failure to comply with instructions. Finally, two participants were excluded after they reported difficulty understanding the questions in a post-study questionnaire. The final sample had 133 participants.

We also excluded a number of propositions that proved ambiguous. Our intention in creating propositions using the most frequently listed object features from the McRae et al. (2005) norms was to have propositions that



were clearly true or false. Unfortunately this method unintentionally resulted in some questions being ambiguous. Consider two such ambiguous propositions: “Does it eat → strawberry” [no]; “Is it hunted by people? → rabbit” [yes]. Strawberries certainly do not eat in the same sense that animals do, but it is conceivable that some people invoke a broader meaning of “eating” that denotes metabolic processes more generally or thought the proposition asked about people eating strawberries. Rabbits *can* be hunted, but are not typical of animals hunted by people.

We identified ambiguous propositions in two ways. First, ambiguous propositions were identified by an independent group of participants who rated the propositions on a 5-point scale from (1) *Definitely No* to (5) *Definitely Yes*. Propositions whose mean truth value did not differ significantly from the midpoint of the scale (*Maybe*) were labeled ambiguous, and excluded from the final sample. Ambiguous propositions were also identified by investigating error rates on the baseline trials without any visual interference in the experiment. Propositions that were not reliably answered correctly without interference were excluded.

Most of the ambiguous propositions were ones for which the correct response was “no”. These false propositions were formed by randomly assigning a cue and feature combination that in some cases resulted in an ambiguous question, e.g., “It is juicy? → bottle”. Others were ambiguous because they were derived from the feature production norms even if they weren’t factually true, e.g., “Does it live in the Arctic? → penguin”.

A total of 139 propositions (26%) were deemed ambiguous by these metrics, and excluded from the final sample. This unintentional limitation was remedied in the preregistered replication in which none of the ambiguous propositions were tested. A plot of effects with and without the ambiguous propositions across all experiment runs is shown in Fig. 3D.

### Statistical modeling

We tested the effect of visual interference on performance using logistic and linear multi-level regression models (lme4 package; Bates et al., 2014) in R (R Core Team, 2014). In all models we allowed overall performance (accuracy, reaction time) to vary by participant (random intercepts).

Our primary measure was the effect of visual interference on the likelihood of making an error. For these analyses, we report the parameter estimates in the units of the logistic regression models (log-odds). Log-odds can be converted to odds via exponentiation. For example, a difference of 1.38 log-odds converts ( $e^{1.38}$ ) to a difference of 4:1 odds in favor of the event, in this case, an increase in the likelihood of making an error. For ease of interpretation, we also report error rates.

For modeling reaction times (ms), we estimated the degrees of freedom for the multi-level linear regression models (and subsequent significance tests) using the `lmerTest` package (Kuznetsova et al., 2014).

### Results

All materials, data, and analyses are available from the Open Science Framework web page for this work <[osf.io/s64ik](https://osf.io/s64ik)>.

### Accuracy analyses

The results of Experiment 2 for all runs of the experiment are shown in Fig. 3. In the original sample, participants were just as likely to make an error on visual questions ( $M = 4.9\%$ ) as they were on nonvisual questions ( $M = 4.5\%$ ), 0.09 log-odds, 95% CI  $[-0.33, 0.51]$ ,  $z = 0.43$ ,  $p = 0.66$ . Visual interference significantly increased errors to visual questions ( $M = 8.7\%$ ), 0.62 log-odds, 95% CI  $[0.25, 1.00]$ ,  $z = 3.2$ ,  $p = 0.001$ , but not to nonvisual questions ( $M = 3.8\%$ ),  $-0.17$  log-odds, 95% CI  $[-0.61, 0.27]$ ,  $z = -0.75$ ,  $p = 0.45$ , resulting in a reliable *interference* × *question type* interaction, 0.79 log-odds, 95% CI  $[0.21, 1.37]$ ,  $z = 2.7$ ,  $p = 0.0069$  (Fig. 3C: **Original**). Including the ambiguous propositions did not change this interaction, 0.47 log-odds, 95% CI  $[0.07, 0.86]$ ,  $z = 2.3$ ,  $p = 0.020$ .

When we conducted the same experiment again one year later, we failed to replicate our original findings,  $-0.32$  log-odds, 95% CI  $[-0.76, 0.11]$ ,  $z = -1.5$ ,  $p = 0.14$ , (Fig. 3C: **Replication**), prompting a careful review of our methods. We identified two potential problems with response biases in the original design (see **Preregistered replication**, above) and also removed ambiguous propositions from the set of all possible questions. We believed these slight modifications would improve the reliability of the effect without changing our primary hypothesis. As expected, in the preregistered replication, the *interference* × *question type* interaction was once again reliable and in the expected direction, 0.59 log-odds, 95% CI  $[0.11, 1.07]$ ,  $z = 2.4$ ,  $p = 0.016$  (Fig. 3C: **Preregistered**). Visual interference led to an increase in error rates when answering visual questions,  $M_{\text{no interference}} = 3.9\%$ ,  $M_{\text{with interference}} = 6.2\%$ , but not nonvisual questions,  $M_{\text{no interference}} = 4.9\%$ ,  $M_{\text{with interference}} = 4.3\%$ .

Combining the original experiment and the preregistered replication shows a highly reliable *interference* × *question type* interaction: 0.67 log-odds, 95% CI  $[0.30, 1.04]$ ,  $z = 3.6051$ ,  $p = 0.0003$ . Collapsing across all three runs of the experiment generally supports the main conclusions drawn from the original sample and the preregistered replication (Fig. 3B). Participants were just as likely to make an error on visual questions as they were on nonvisual questions,  $M_{\text{visual}} = 4.1\%$ ,  $M_{\text{nonvisual}} = 5.5\%$ , 0.14 log-odds, 95% CI  $[-0.06, 0.35]$ ,  $z = 1.4$ ,  $p = 0.16$ . With visual interference, errors became significantly higher on visual questions ( $M = 5.5\%$ ), 0.30 log-odds, 95% CI  $[0.11, 0.49]$ ,  $z = 3.2$ ,  $p = 0.0014$ , but not on nonvisual questions ( $M = 3.7\%$ ), 0.04 log-odds, 95% CI  $[-0.17, 0.25]$ ,  $z = 0.42$ ,  $p = 0.68$ , resulting in a marginal *interference* × *question type* interaction, 0.25 log-odds, 95% CI  $[-0.03, 0.53]$ ,  $z = 1.8$ ,  $p = 0.072$ .<sup>2</sup>

<sup>2</sup> The impact of the ambiguous propositions is shown in Fig. 3D. Overall the *interference* × *question type* interaction was stronger when ambiguous propositions were excluded, 0.25 log-odds, 95% CI  $[-0.03, 0.53]$ ,  $z = 1.8$ ,  $p = 0.072$ , than when they were included, 0.14 log-odds, 95% CI  $[-0.06, 0.34]$ ,  $z = 1.4$ ,  $p = 0.17$ . Given the difficulty with evaluating correctness for the ambiguous propositions, we excluded these propositions in the preregistered replication attempt. Without any ambiguous propositions, visual interference affects performance on visual questions only, 0.59 log-odds, 95% CI  $[0.11, 1.07]$ ,  $z = 2.4$ ,  $p = 0.016$ .

### Reaction times

We did not find any selective effect of visual interference on RTs. Unexpectedly, visual interference resulted in faster response times for *both* visual and nonvisual questions. This was true both for the original design when participants were prevented from responding until the offset of the auditory cue,  $-13$  ms, 95% CI  $[-22, -5]$ ,  $z = -3.2$ ,  $p = 0.0016$ , and in the preregistered design where the start of the response window was simultaneous with cue onset,  $-33$  ms, 95% CI  $[-43, 22]$ ,  $z = -6.0$ ,  $p < -0.001$ .<sup>3</sup> This main effect of visual interference did not vary for visual or nonvisual questions, a difference of only 1 ms, 95% CI  $[-12, 14]$ ,  $z = 0.13$ ,  $p = 0.89$ .

Although there were no baseline differences in accuracy or rated difficulty between visual and nonvisual questions, visual questions took 25 ms longer to answer correctly than nonvisual questions, 95% CI  $[18, 32]$ ,  $z = 7.4$ ,  $p < 0.001$ , raising the worry that visual questions were somehow more difficult, and the mask made them more difficult still. We address this possibility in several ways in the sections below.

### Ruling out confounds

In addition to selecting visual and nonvisual properties that were equated on production frequency (see **Materials**, above), we also explored potential differences between visual and nonvisual properties based on our own norming procedures.

*Was the effect of visual interference confounded by question difficulty?* An alternative interpretation of our results is that visual questions are more difficult than nonvisual questions, a difference that was increased by visual interference. Several results speak against this interpretation. First, the effect of visual interference was not greater for more difficult questions as rated by naïve participants in the norming study, 0.022 log-odds, 95% CI  $[-0.14, 0.18]$ ,  $z = 0.28$ ,  $p = 0.78$ . A different measure of difficulty was obtained by measuring baseline performance (RTs) on trials without visual interference. This measure of baseline difficulty was not correlated with the difficulty as rated by the norming participants,  $r = -0.039$ , n.s., or with accuracy on no-interference trials,  $r = 0.05$ , n.s. As with the subjective difficulty ratings, the effect of visual interference did not vary by baseline RTs to the question, 0.0007 log-odds,  $z = 1.5$ ,  $p = 0.13$ . A third measure of difficulty as measured by error rates on trials without visual interference was positively correlated with normed difficulty,  $r = 0.27$ , and by this measure, questions that were more difficult were likely to be more affected by visual interference,  $-4.77$  log-odds, 95% CI  $[-0.14, 0.18]$ ,  $z = -5.0$ ,  $p < 0.001$ . However, including baseline error rate as a covariate did not change the question type-by-visual interference interaction, 0.30 log-odds,  $z = 2.0$ ,  $p = 0.041$ . Including baseline

RT as a covariate did not change the direction of the effect, 0.25 log-odds,  $z = 1.80$ ,  $p = 0.07$ .

*Was the effect of visual interference confounded by question abstractness.* To further rule out the possibility that visual interference disrupted performance as a function of something other than visualness, we explored whether the effect of the interference was moderated by question abstractness. In the analyses reported above, we distinguished between visual properties and other properties, but collapsed between nonvisual abstract knowledge (e.g., “Is a donkey a mammal?”) and knowledge concerning other modalities (“Does a donkey make noise?”). If visual questions were less abstract and that is why they were more affected by visual interference, then the effect of visual interference should be greater for less abstract questions regardless of sensory modality. Visual questions were indeed viewed as requiring more sensory knowledge/less abstract knowledge ( $M = 1.9$ ) compared to nonvisual questions ( $M = 2.4$ ),  $p < 0.001$ . However, the effect of visual interference was not sensitive to the abstractness of the question, 0.015 log-odds, 95% CI  $[-0.13, 0.16]$ ,  $z = 0.21$ ,  $p = 0.83$ .

### Discussion

Experiment 2 shows that people’s ability to make use of their visual knowledge—long-term memory of what things look like—is disrupted by a simple form of visual interference. Visual interference disrupted performance strictly as a function of visual knowledge rather than question difficulty, truth-value, or question abstractness. These results highlight the selectivity of the interference: viewing colorful flashing rectangles affected visual knowledge, but not encyclopedic and nonvisual sensory knowledge. The most parsimonious explanation of the finding that people’s ability to answer simple questions about what things look like can be disrupted by irrelevant visual inputs is that visual knowledge is represented in a perceptual format. We discuss alternative possibilities including the role of visual imagery and of retrieval processes in the General Discussion.

### General discussion

Simply viewing colorful rectangles interfered with people’s understanding of the visual aspects of verbally cued categories. Answering yes/no questions about visual properties (Fig. 3) and judging which object is correctly oriented (Fig. 1) are tasks that unequivocally require visual knowledge. That visual interference reduced the accuracy in answering questions requiring visual knowledge and reduced the verbal cueing effect strongly suggests that making use of visual knowledge is mediated by perceptual mechanisms. These results support theories that posit that conceptual representations are distributed such that visual attributes are represented in a perceptual format (Barsalou et al., 2003; Pulvermüller, 2005), or at least a format that can be easily interfered with by perceptual inputs. Our results are difficult to accommodate by views positing that

<sup>3</sup> Ironically part of the motivation for modifying the response window in the preregistered design was to remedy what we found to be a confusing improvement in correct response RTs by visual interference in the original design. Instead, the paradoxical facilitation of RTs by the mask became even larger. This effect may be partly explained by participants being pressured to respond more quickly by the rapidly flashing masks, making accuracy the better dependent variable.

visual knowledge is represented in an amodal format (Mahon & Caramazza, 2008).

We recognize, of course, that people's ability to answer even the most visual questions remained at a high level (90–95%) on the visual interference trials (Fig. 3B). Given how superficial our visual interference task is—we are simply asking people to look at flashing rectangles—it should not be surprising that people can still answer these questions quite accurately. The critical point is that even this very simple and superficial form of visual interference resulted in a selective increase in errors, strictly as a function of how much visual knowledge was required (Fig. 3). The crux of our results therefore lies not with the absolute size of the accuracy or RT differences, but with demonstrating that visual interference disrupts visual, and specifically visual, knowledge. We believe this finding is difficult, perhaps impossible to accommodate without positing that visual knowledge is at least partly represented using a visual format.

It may be argued that our effects depend on participants utilizing an explicit mental imagery strategy, and what the visual interference is doing is disrupting these mental images. We believe that the fast-paced multi-trial nature of the experiments makes an *explicit* mental imagery strategy—e.g., of the type studied by Kosslyn et al. (2001) and Farah (1989)—unlikely, although we cannot rule it out. If such “imagery” is causally implicated in visual knowledge, it is not clear to us that this explanation is a true alternative to our proposal, as it would mean that answering simple questions about visual properties, and to reap the benefits of verbal cueing, depends on visual imagery (a perceptual process). Rather than attempting to separate the visual aspects of object knowledge from mental imagery—implicit or explicit—it may be more accurate to consider mental imagery, visual long-term memory, and also visual working memory all as processes that utilize mechanisms shared with visual perception (Albers et al., 2013; Naselaris et al., 2015; Pearson, Clifford, & Tong, 2008).

A question left open by our work is whether visual knowledge is *stored* within the perceptual system (e.g., extrastriate cortex) or if the visual system, broadly construed is only involved in accessing such knowledge, which is itself stored elsewhere. It is unclear what kind of evidence could distinguish these accounts and whether such a distinction is sensible given the difficulty of drawing strict boundaries between content and process in the brain (Anderson, 2011; Pobric, Jefferies, & Lambon Ralph, 2007; Pobric, Jefferies, & Lambon Ralph, 2010; Price & Devlin, 2003).

We do not wish to argue that all knowledge is dependent on modality-specific mechanisms. Knowledge of non-visual “encyclopedic” properties was not disrupted by visual interference and may well be represented in a more abstract non-perceptual format. The extent to which modality congruent interference disrupts knowledge specific to that modality (e.g., auditory interference affecting auditory knowledge) is left open for future work. How we learn something may also prove important. For example, that alligators have sharp teeth may be inferred after learning that alligators are carnivores without ever having seen an alligator. A key question for future research is how

different learning experiences shape how knowledge is represented across individuals.

Although our work examined only the visual modality, we expect that knowledge about other modalities similarly depends on perceptual mechanisms and is also subject to modality-appropriate interference. For example auditory interference should affect auditory knowledge (Vallet, Brunel, & Versace, 2010). That some domain-specific knowledge depends on perceptual states is also supported by findings that knowledge of actions and movement depends in part on motor resources (Pulvermüller, 2005; Yee, Chryssikou, Hoffman, & Thompson-Schill, 2013), and that visual or auditory memory load can disrupt property verification in the loaded modality (Vermeulen, Corneille, & Niedenthal, 2008).

## Conclusion

Simply viewing colorful patterns disrupted behaviors that required visual knowledge, a result that is parsimoniously explained by the view that visual perception and visual knowledge rely on common mechanisms. Although much of human knowledge abstracts away from perceptual details, knowledge of what things look like appears to be represented in a visual format.

## Acknowledgments

The study was partially funded by NSF-PAC #1331293 to G.L. We thank Callie Porter-Borden, Will Reinhart, Oliver Roe, Nita Sharma, Heather Wolf, Nita Sharma, Daniel Zych, and Jesse Reid for helping collect the data.

## Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jml.2016.07.002>.

## References

- Albers, A. M., Kok, P., Toni, I., Dijkerman, H. C., & de Lange, F. P. (2013). Shared representations for working memory and mental imagery in early visual cortex. *Current Biology*, 23(15), 1427–1431. <http://dx.doi.org/10.1016/j.cub.2013.05.065>.
- Amsel, B. D., Urbach, T. P., & Kutas, M. (2014). Empirically grounding grounded cognition: The case of color. *NeuroImage*, 99C, 149–157. <http://dx.doi.org/10.1016/j.neuroimage.2014.05.025>.
- Anderson, B. (2011). There is no such thing as attention. *Frontiers in Psychology*, 2. <http://dx.doi.org/10.3389/fpsyg.2011.00246>.
- Barsalou, L. W., Simmons, W. K., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*, 7(2), 84–91. [http://dx.doi.org/10.1016/S1364-6613\(02\)00029-3](http://dx.doi.org/10.1016/S1364-6613(02)00029-3).
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). lme4: Linear mixed-effects models using Eigen and S4. R package version 1.1-7 Retrieved from <http://CRAN.R-project.org/package=lme4>.
- Boutonnet, B., & Lupyan, G. (2015). Words jump-start vision: A label advantage in object recognition. *The Journal of Neuroscience*, 32(25), 9329–9335.
- Caramazza, A., Hillis, A. E., Rapp, B. C., & Romani, C. (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology*, 7(3), 161–189. <http://dx.doi.org/10.1080/02643299008253441>.

- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407–428. <http://dx.doi.org/10.1037/0033-295X.82.6.407>.
- Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, 132(2), 163–201. <http://dx.doi.org/10.1037/0096-3445.132.2.163>.
- Farah, M. J. (1989). The neural basis of mental imagery. *Trends in Neurosciences*, 12(10), 395–399. [http://dx.doi.org/10.1016/0166-2236\(89\)90079-9](http://dx.doi.org/10.1016/0166-2236(89)90079-9).
- Firestone, C., & Scholl, B. J. (2015). Enhanced visual awareness for morality and pajamas? Perception vs. memory in “top-down” effects. *Cognition*, 149, 409–416.
- Hoffman, P., & Lambon Ralph, M. A. (2013). Shapes, scents and sounds: Quantifying the full multi-sensory basis of conceptual knowledge. *Neuropsychologia*, 51(1), 14–25. <http://dx.doi.org/10.1016/j.neuropsychologia.2012.11.009>.
- Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience*, 2(9), 635–642. <http://dx.doi.org/10.1038/35090055>.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2014). lmerTest: Tests in linear mixed effects models Retrieved from <<http://CRAN.R-project.org/package=lmerTest>>.
- Lupyan, G., & Spivey, M. J. (2010). Redundant spoken labels facilitate perception of multiple items. *Attention, Perception, & Psychophysics*, 72(8), 2236–2253. <http://dx.doi.org/10.3758/APP.72.8.2236>.
- Lupyan, G., & Thompson-Schill, S. L. (2012). The evocative power of words: Activation of concepts by verbal and nonverbal means. *Journal of Experimental Psychology: General*, 141(1), 170–186. <http://dx.doi.org/10.1037/a0024904>.
- Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences*, 110(35), 14196–14201. <http://dx.doi.org/10.1073/pnas.1303312110>.
- Mahon, B. Z., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology-Paris*, 102(1–3), 59–70. <http://dx.doi.org/10.1016/j.jphysparis.2008.03.004>.
- Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology*, 58, 25–45. <http://dx.doi.org/10.1146/annurev.psych.57.102904.190143>.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37(4), 547–559. <http://dx.doi.org/10.3758/BF03192726>.
- Meteyard, L., Zokaei, N., Bahrami, B., & Vigliocco, G. (2008). Visual motion interferes with lexical decision on motion words. *Current Biology*, 18(17), R732–R733. <http://dx.doi.org/10.1016/j.cub.2008.07.016>.
- Naselaris, T., Olman, C. A., Stansbury, D. E., Ugurbil, K., & Gallant, J. L. (2015). A voxel-wise encoding model for early visual areas decodes mental images of remembered scenes. *NeuroImage*, 105, 215–228. <http://dx.doi.org/10.1016/j.neuroimage.2014.10.018>.
- Papeo, L., Pascual-Leone, A., & Caramazza, A. (2013). Disrupting the brain to validate hypotheses on the neurobiology of language. *Frontiers in Human Neuroscience*, 7, 148. <http://dx.doi.org/10.3389/fnhum.2013.00148>.
- Pearson, J., Clifford, C. W. G., & Tong, F. (2008). The functional impact of mental imagery on conscious perception. *Current Biology*, 18(13), 982–986. <http://dx.doi.org/10.1016/j.cub.2008.05.048>.
- Pinker, S. (1994). *The language instinct*. New York: W. Morrow and Co.
- Pobric, G., Jefferies, E., & Lambon Ralph, M. A. (2007). Anterior temporal lobes mediate semantic representation: Mimicking semantic dementia by using rTMS in normal participants. *Proceedings of the National Academy of Sciences*, 104(50), 20137–20141. <http://dx.doi.org/10.1073/pnas.0707383104>.
- Pobric, G., Jefferies, E., & Lambon Ralph, M. A. (2010). Category-specific versus category-general semantic impairment induced by transcranial magnetic stimulation. *Current Biology*, 20(10), 964–968. <http://dx.doi.org/10.1016/j.cub.2010.03.070>.
- Price, C. J., & Devlin, J. T. (2003). The myth of the visual word form area. *NeuroImage*, 19(3), 473–481. [http://dx.doi.org/10.1016/S1053-8119\(03\)00084-3](http://dx.doi.org/10.1016/S1053-8119(03)00084-3).
- Pulvermüller, F. (2005). Brain mechanisms linking language and action. *Nature Reviews Neuroscience*, 6(7), 576–582. <http://dx.doi.org/10.1038/nrn1706>.
- Pylshyn, Z. W. (1986). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge: MIT Press.
- R Core Team (2014). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <<http://www.R-project.org/>>.
- Rossion, B., & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object pictorial set: The role of surface detail in basic-level object recognition. *Perception*, 33(2), 217–236.
- Simmons, W. K., Ramjee, V., Beauchamp, M. S., McRae, K., Martin, A., & Barsalou, L. W. (2007). A common neural substrate for perceiving and knowing about color. *Neuropsychologia*, 45(12), 2802–2810. <http://dx.doi.org/10.1016/j.neuropsychologia.2007.05.002>.
- Tulving, E. (1972). Episodic and semantic memory. In *Organization of memory*. Oxford, England: Academic Press, pp. xiii, 423.
- Vallet, G., Brunel, L., & Versace, R. (2010). The perceptual nature of the cross-modal priming effect: Arguments in favor of a sensory-based conception of memory. *Experimental Psychology*, 57(5), 376–382. <http://dx.doi.org/10.1027/1618-3169/a000045>.
- Vermeulen, N., Corneille, O., & Niedenthal, P. M. (2008). Sensory load incurs conceptual processing costs. *Cognition*, 109(2), 287–294. <http://dx.doi.org/10.1016/j.cognition.2008.09.004>.
- Yee, E., Chrysikou, E. G., Hoffman, E., & Thompson-Schill, S. L. (2013). Manual experience shapes object representations. *Psychological Science*, 24(6), 909–919. <http://dx.doi.org/10.1177/0956797612464658>.