Where word and world meet: Language and vision share an abstract representation of symmetry^{*}

Alon Hafri^{a,b,1}, Lila R. Gleitman^c, Barbara Landau^b, & John C. Trueswell^c

^aDepartment of Psychological & Brain Sciences, Johns Hopkins University, Baltimore, MD 21218

^bDepartment of Cognitive Science, Johns Hopkins University, Baltimore, MD 21218

^cDepartment of Psychology, University of Pennsylvania, Philadelphia, PA 19104

Abstract

Symmetry is ubiquitous in nature, in logic and mathematics, and in perception, language, and thought. Although humans are exquisitely sensitive to visual symmetry (e.g., of a butterfly), symmetry in natural language goes beyond visuospatial properties: many words point to abstract concepts with symmetrical content (e.g., *equal*, *marry*). For example, if Mark marries Bill, then Bill marries Mark. In both cases (vision and language), symmetry may be formally characterized as invariance under transformation. Is this a coincidence, or is there some deeper psychological resemblance? Here we asked whether representations of symmetry correspond across language and vision. To do so, we developed a novel cross-modal matching paradigm. On each trial, participants observed a visual stimulus (either symmetrical or non-symmetrical) and had to choose between a symmetrical and non-symmetrical English predicate unrelated to the stimulus (e.g., “negotiate” vs. “propose”). In a first study with visual events (symmetrical collision or asymmetrical launch), participants reliably chose the predicate matching the event’s symmetry. A second study showed that this “language-vision correspondence” generalized to objects, and was weakened when the stimuli’s binary nature was made less apparent (i.e., for one object, rather than two inward-facing objects). A final study showed the same effect when nonsigners guessed English translations of signs from American Sign Language, which expresses many symmetrical concepts spatially. Taken together, our findings support the existence of an abstract representation of symmetry which humans access via both perceptual and linguistic means. More broadly, this work sheds light on the rich, structured nature of the language-cognition interface.

Keywords: cross-modal; abstract language; visual relations; language-cognition interface; conceptual structure

^{*} All data, code, analyses, stimuli, and pre-registrations are available from our project’s Open Science Framework repository at this link: <https://osf.io/64dnq>. Hypotheses, sample sizes, and analysis plans (as well as other details) for Experiments 1a, 2, and 3 were pre-registered. Preliminary analyses of Experiments 1a, 1b, and 2 originally appeared in conference proceedings (Hafri et al., 2021).

¹ To whom correspondence should be addressed.

Email address: alon@udel.edu

Introduction

Look at the images in Figure 1a. Beyond the variety of content that is depicted, it is clear that they have something in common: they are all symmetrical. Symmetry is pervasive in both the natural and constructed world — in biological and physical systems (as in the structure of plants, animals, or crystals), artistry (as in sculptures or paintings), architecture (as in the Eiffel Tower or the Coliseum), and more (Weyl, 1952). And the human visual system is exquisitely sensitive to the symmetrical structure of images like those in Figure 1a (for a review, see Wagemans, 1997).

Yet symmetry goes far beyond the visual and sensory world; it is also present in language and cognition, which represent concepts in ways that also invoke symmetry. Consider Figure 1b. If Mark and Bill marry, we know that Mark marries Bill and Bill marries Mark; there is a symmetry inherent in the event being described. Such symmetry is threaded throughout the vocabulary of natural languages, with virtually every part of speech having words for both symmetrical and non-symmetrical concepts: nouns (*cousin* vs. *father*), verbs (*marry* vs. *adopt*), adjectives (*similar* vs. *larger*), prepositions (*near* vs. *above*), and conjunctions (*and* vs. *because*). As such, symmetry has been a topic of great interest to linguists and psycholinguists, who have attempted to characterize the syntactic behavior of symmetrical terms across languages, and to understand how such terms are acquired (e.g., Chestnut & Markman, 2016; Dimitriadis, 2008; Gleitman et al., 1996; Gleitman & Partee, in review; Miller, 1998; Partee, 2008; Siloni, 2012; Winter, 2018).

Figure 1. Symmetry in vision and language. (a) Symmetry appears throughout the natural and constructed world, and is especially visually salient for left-right mirror-reflective symmetry. (b) Symmetry is also threaded throughout the vocabulary of natural language: many lexical items like *marry*

have symmetrical entailments associated with them. In both cases, (a) and (b), symmetry can be characterized as invariance under transformation. (Photo credits in (a) are the following. Top-left: Faruk Kaymak at https://unsplash.com/photos/P_Ne56WEe5s; Top-right: Erol Ahmed at <https://unsplash.com/photos/aIYFR0vbADk>; Bottom: <https://www.rawpixel.com/image/6043543>.)

More formally, a relation R is symmetrical if and only if for all x, y : if $R(x, y)$, then $R(y, x)$ (Partee, 2008; Winter, 2018). This symmetrical entailment puts symmetry on par with other highly abstract logical concepts, including negation (e.g., *no*, *not*, *never*) and quantification (e.g., *one*, *some*, *all*). Such properties and their combinatorics are essential to how we construct, understand, and evaluate utterances. For example, words for negation invert the truth-value of propositions (contrast *It's raining* with *It's not raining*), and quantifiers pick out sets (contrast *She has all the flowers* with *She has some of the flowers*). Likewise, symmetrical relations underlie much of logical and scientific reasoning, as when we reason whether two quantities or entities are similar or even equal — a difficulty that many children in their early school years struggle to overcome (e.g., Johannes & Davenport, 2017; McNeil, 2014; National Council of Teachers of Mathematics, 2000).

Symmetry Across Language and Vision?

Perhaps the appearance of symmetry in both vision and language is a mere coincidence, but it may also be that they share a deeper psychological commonality. Thus, investigating the mental representation of symmetry presents a unique opportunity to ask a fundamental psychological question: to what degree do abstract properties correspond across different cognitive systems? Specifically, we ask the following: (1) Is the abstract logical symmetry of linguistic terms mentally accessible, even when such terms are presented in isolation? (2) Is the format of symmetry that is accessed via visual perception also in some cases abstract and relational? and (3) Do linguistic and visual representations of symmetry correspond with one another in the mind, and under what conditions?

Beyond revealing how representations of symmetry correspond across different cognitive systems, answering these questions will also shed light on how children might learn symmetrical terms, and how adults represent them. In particular, the abstract and even lawlike nature of symmetrical meanings presents a problem for the language learner, who must discover which unfamiliar words map onto abstract meanings (Fisher et al., 2019; Gleitman, 1990; Gleitman et al., 2019; Landau & Gleitman, 1985) — a process which, especially for symmetrical concepts, requires projecting from a finite set of instances to a more general and lawlike property (Goodman, 1955). A correspondence between linguistic and visual representations of symmetry may thus be part of the toolkit the learner uses to solve this problem (a possibility we outline in more detail in the General Discussion).

Distinct Symmetries in Language and in Vision

The existence of visual symmetry is not controversial. In vision, figural symmetry of single objects or spatial arrays is extracted rapidly and automatically, and functions as a Gestalt property of perceptual organization (for reviews, see Wagemans, 1997; Wagemans et al., 2012). It is also available early in infancy, arriving by four months of age (Bornstein & Krinsky, 1985). Mirror-reflective symmetry — especially vertical left-right symmetry — is particularly salient (relative to other types, such as rotational and translational symmetry, or symmetry along the horizontal or diagonal axes; Palmer & Hemenway, 1978). Intriguingly, symmetry

representation in perception is not limited to the visual modality, but appears to be a general spatial ability: it is also observed in the tactile modality in the congenitally blind (Cattaneo et al., 2010).

In contrast to visual symmetry, the logical symmetry of lexical items is less transparent, as it is not overtly marked morphologically in English. Yet there is a wealth of evidence that such symmetry exists, and that it has many interesting properties (Dimitriadis, 2008; Gleitman & Partee, in review; Partee, 2008; Siloni, 2012) — including problematic ones for some aspects of linguistic theory (e.g., that each noun phrase of a predicate must be assigned just one unique thematic role; Chomsky, 1981).

Such symmetry can be appreciated by examining the sentences below:

- (1) Mark and Bill marry.
- (2) Mark and Bill marry one another.
- (3) Mark marries Bill.
- (4) Bill marries Mark.
- (5) ?Mark marries.

It can be observed that if (1) is true, it entails the reciprocal (2), as well as the sub-event in (3) and this sub-event's inverse (4). Notice also that because symmetry is a property of *binary* relations, an intransitive utterance with a singular subject (5) is semantically anomalous. These properties — that the intransitive must have a semantically plural subject and that it has a roughly reciprocal interpretation — have been identified as a key two-part “litmus test” for whether a lexical item has a symmetrical meaning (Gleitman et al., 1996; see also Gleitman & Partee, in review).

It is also important to note that symmetry goes beyond mere collectivity, which describes a situation in which two or more entities participate together in some relation. Many linguistic terms can indicate a collective event or situation; for example, just as *Mark and Bill marry* entails both that *Mark got married* and *Bill got married*, likewise *Mark and Bill drown* entails both that *Mark drowns* and that *Bill drowns*. However, symmetry entails a reciprocity or mutuality that goes beyond mere collectivity: *Mark and Bill marry* entails that they marry each other, but *Mark and Bill drown* does not entail that they drown each other. Although there has been some debate about whether one particular predicate, *similar*, is truly a symmetrical concept (Tversky, 1977), Gleitman and colleagues (1996) convincingly demonstrated that *similar* and similar concepts are logically symmetrical, with apparent “asymmetrical” properties introduced by particular syntactic structures (e.g., asymmetries in subject and complement positions, such as *Mark is similar to Bill*) (see also Chestnut & Markman, 2016; Landau & Gleitman, 2015).

Invariance Under Transformation: A Common Symmetry?

Although symmetry appears in both vision and language, it is unknown whether symmetry is represented in ways that invite mental correspondences across these domains. After all, there are many properties that are distinctly visual or distinctly linguistic (e.g., the property of material texture in vision, or the type/token distinction in language, to name a few), for which there are unlikely to be obvious mental correspondences across domains (for a detailed discussion, see Jackendoff, 1987). Is it just a coincidence, two areas of cognitive science employing the same term (“symmetry”) to describe two different properties?

There are important reasons to think not. Consider again Figure 1. Notice that in both cases —visual and linguistic — the symmetry can be characterized as *invariance under transformation*. In the spatial domain, symmetry is a property of figures or patterns that are invariant under transformations such as rotation, reflection, or translation (Weyl, 1952). For example, a butterfly's wings show reflective symmetry, such that if

one half of the butterfly image in Figure 1a were “copied over” to the other half, the resulting image content would remain the same. In the logical domain, this invariance is one in which the arguments of a binary relation are inverted while maintaining roughly the same truth conditions: a relation R is symmetrical if and only if for all x, y : if $R(x, y)$ is true, then $R(y, x)$ must also be true (Partee, 2008). Thus, if x equals y , then y must equal x . Linguistic predicates such as *marry* token relational concepts with these entailments.²

Although it is conceivable that invariance under transformation might underlie the mental representation of symmetry in both cases, it might be that this is merely an abstract theoretical description of their similarity — one to which the mind does not have access. After all, consider how different the units are over which these symmetries operate. In the visuospatial domain, symmetry holds over visuospatial features like textures, contours, or shapes. In the logical domain, symmetry holds for relational concepts whose arguments are also concepts (of entities). Can invariance under transformation rise above these differences, such that the mind treats both forms of symmetry as similar?

The Present Studies: Cross-Modal Matching

Here we asked whether mental representations of symmetry correspond across language and vision, and we characterize the conditions under which this correspondence is available to the mind. To do so, we developed a novel cross-modal matching paradigm, based on a method that Strickland and colleagues (2015) recently used to reveal a bias to map linguistic telicity or “boundedness” (in English predicates) to visual boundedness (in signs from different sign languages). This paradigm bears similarities to match-to-sample tasks that have been used for investigating mappings within modality (e.g., Hochmann et al., 2017; Shao & Gentner, 2019).³

On each trial of our experiments, participants viewed a visual stimulus (symmetrical or non-symmetrical), and then had to choose between two English predicates (one symmetrical, one non-symmetrical) that best matched what they saw (Figure 2). Crucially, the predicates were not semantically related to the visual stimuli in any direct way, apart from the property of symmetry. We predicted that participants would prefer to “match” symmetrical predicates with symmetrical visual stimuli, and non-symmetrical predicates with non-

² We are simplifying the situation a bit here, as certain symmetrical predicates (e.g., *kiss*, *hug*) may not indicate a reciprocated event in the transitive — apparently violating the mutual entailment of a symmetrical relation (i.e., if Mark kisses Bill it does not necessarily entail that Bill kisses Mark). We sidestep this complication in our studies by only using terms that meet a stricter logical definition of symmetry: e.g., *marry* or *combine*, in which even the transitive forms entail their inverse (e.g., Mark marries Bill implies that Bill marries Mark). However, we note that even apparent “problem cases” like *kiss* and *hug* have a reciprocal interpretation when in the intransitive (Mark and Bill kiss) and so obey the “litmus test” discussed above. This suggests that the asymmetries for these cases are introduced largely by the asymmetrical syntactic structures (Gleitman et al., 1996; Landau & Gleitman, 2015). See Gleitman & Partee (in review) for a thorough discussion of these issues, and an attempt to provide classifications for these different types of symmetricals.

³ Within-modality mapping tasks are often used in studies on analogical reasoning or structural alignment processes, for which the goal is often to explore how humans make metaphors, analogies, or similarity judgments across stimuli (e.g., Gattis, 2004; Goldwater & Gentner, 2015; Goldwater et al., 2011; Markman & Gentner, 1993), and how this ability develops (e.g., Shao & Gentner, 2019). In investigating these processes, such work generally *presupposes* that the compared representations have certain kinds of entities, attributes, and even relations in common (Markman & Gentner, 2000). By contrast, our goal here was to provide evidence for the existence of a particular relational property itself — namely, symmetry — and to investigate its nature, by exploiting the fact that our task required comparison between visual stimuli and spoken-language stimuli. Some readers might interpret this comparison problem as one that demands that the observer carry out an analogy. However, our view is different, as we elaborate on in the General Discussion (see the section “Possible Mechanisms of Cross-Modal Matching”), where we also discuss alternative accounts that do not invoke conceptual representations.

symmetrical visual stimuli; in other words, we predicted that this association would be intuitive, such that out of all the many possible ways words and visual forms may be matched together, the dimension of symmetry would be especially prominent in the mind.

In our cross-modal matching studies, we paired various kinds of visual stimuli with the same set of English predicate pairs to ask what aspects of visual symmetry invite a mapping to the abstract symmetry of linguistic terms. In Experiments 1a and 2, we used simple stimuli with geometric shapes to isolate specific aspects of visual symmetry. In particular, in Experiment 1a, we used dynamic visual events (symmetrical collision vs. asymmetrical launch) of the kind known to elicit robust causal representations (Michotte 1946/1963), and in Experiment 1b, we collected symmetry ratings for the predicates to ensure that the effects in Experiment 1a were attributable to symmetry per se. In Experiment 2, we used static objects, manipulating the number of objects (two vs. one) to ask whether binary relationships between entities are crucial for the interpretation of a symmetrical relation in vision. Finally, in Experiment 3, we used American Sign Language (ASL) signs that were either visually symmetrical or non-symmetrical (with none of the ASL signs being translations of the English predicate choices on a trial and with none of the participants being familiar with ASL). To preview our results, we found a robust matching effect across experiments, and we determined that the binary nature of the stimulus is important to engender a robust correspondence. Readers can experience demos of the tasks for themselves at <https://langdevresearch.org/symmetry>.

Experiment 1a: Visual Events

Do linguistic and visual representations of symmetry correspond with one another in the mind? To address this question, we designed visual and linguistic stimulus sets with some special properties. Here, we exploited a compelling visual phenomenon: causal perception. When observers view such events, they experience the transfer of force rapidly and automatically, and in ways that are phenomenologically compelling and genuinely perceptual (e.g., they show retinotopic visual adaptation; Michotte 1946/1963; Rolfs et al., 2013). Crucially, both asymmetrical causal “launching” and symmetrical “bouncing” events have been studied in visual perception research (Meyerhoff & Suzuki, 2018; Sekuler et al., 1997). To ask whether symmetrical and asymmetrical causal events would be sufficient to engender a cross-domain correspondence, we showed such visual events one at a time to observers and asked them to report which of two predicates (symmetrical or non-symmetrical) “best matched” the visual event they observed.

Figure 2. Example stimuli and trial structure for Experiment 1a. (a) Visual stimuli were symmetrical collisions and non-symmetrical causal launches. (b) On each trial, participants observed a visual stimulus and then selected the phrase that “best matched” the visual display; one was a symmetrical predicate and one was a non-symmetrical predicate. The trial structure was similar for Experiments 2 and 3.

Open Science Practices

All data, code, analyses, stimuli, and pre-registrations (for this experiment and all others reported here) are available at <https://langdevresearch.org/symmetry>. This webpage also includes demos of each experiment, so that readers can experience these tasks as participants did. The sample sizes and analysis plans (as well as other details) for all cross-modal matching experiments were pre-registered.

Participants

Sixty adults (U.S. IP addresses only) were recruited from the online platform Amazon Mechanical Turk (MTurk). (For a discussion of the reliability of this subject pool, see Peer et al., 2017). This sample size was chosen based on pilot data (which are not included in the samples reported here). Sample sizes were pre-registered for this and all other cross-modal matching experiments. All studies were approved by the Institutional Review Boards at Johns Hopkins and the University of Pennsylvania. Participants were paid \$0.80 for their participation, with an average study duration of 6 minutes.

Stimuli

There were two types of stimuli: linguistic (English predicates) and visual (events between geometric shapes).

English Predicates. We generated 24 pairs of predicates, each with one symmetrical and one non-symmetrical item (see Table 1). Predicates ranged from concrete to abstract.

Table 1: Symmetrical (Sym) and Non-Symmetrical (Non-Sym) Predicate Pairs.

	Sym	Non-Sym		Sym	Non-Sym
1.	agree	consent	13.	be equal	exceed
2.	box	punch	14.	be identical	be inferior
3.	chat	tell	15.	interact	intervene
4.	clash	confront	16.	intersect	interfere
5.	collaborate	contribute	17.	marry	adopt
6.	collide	hit	18.	match	gauge
7.	combine	expand	19.	meet	greet
8.	correspond	contact	20.	negotiate	propose
9.	date	befriend	21.	separate	withdraw
10.	debate	lecture	22.	be similar	be typical
11.	differ	alter	23.	tango	lead
12.	disagree	reject	24.	unite	dominate

Symmetrical predicates had to meet several criteria. First, they had to exhibit *logical* symmetry: for all x, y , if $R(x, y)$ then $R(y, x)$. For example, if x equals y then y equals x , but if Mark drowned Bill, it does not entail that Bill drowned Mark. Second, they had to exhibit the *linguistic reflexes* of symmetry (i.e., they had to obey the two-part “litmus test” discussed earlier): (a) their intransitive entails the reciprocal, as in (1) and (2) above, and (b) they are infelicitous with a conceptually singular subject (Gleitman et al., 1996; Gleitman & Partee, in review). If x and y are equal, x and y are equal *to each other*, but the utterance *Mark and Bill drowned* does not entail that Mark and Bill drowned each other. Furthermore, the utterance *x is equal* is infelicitous. These linguistic criteria were used to inform decisions about logical symmetry. For example, if Mark loves Bill, then Bill may in fact love Mark. But the sentence *Mark and Bill love* is infelicitous, and does not entail that they love each other.

We identified symmetrical predicates from several sources: the original list from Gleitman et al. (1996); VerbNet (<http://verbs.colorado.edu/verb-index/index.php>; Kipper et al., 2008), a database of English verbs organized into classes based on their subcategorization frames (an updated form of Levin’s (1993) classes); and from discussions among the authors and colleagues.

To select a non-symmetrical foil for each symmetrical, we found a predicate as close in meaning as we could, apart from the property of symmetry. For example, “to chat” and “to tell” are both verbs of conversing, but *chat* tokens a symmetrical relation, while *tell* does not. We also ensured that non-symmetricals failed the logical and linguistic symmetrical criteria outlined above. Predicate pairs were closely matched in length, log frequency, and concreteness (based on norms in Brysbaert et al., 2014; see Appendix A for details on these

norms). Since in English, collectivity is sometimes marked with the prefix “co-” (e.g., *collaborate*), we also attempted to balance the prevalence of such items across the set of symmetricals and non-symmetricals. Each non-symmetrical item was always “yoked” to the same symmetrical item, such that the same predicates were always presented together as a pair. Predicate pairs appear in Table 1.

Visual Events. In symmetrical events, two objects approached one another, made contact, and bounced away, always with equal velocity (Meyerhoff & Suzuki, 2018; Sekuler et al., 1997). Non-symmetrical events were Michottean causal launches (Michotte, 1946/1963; Rolfs et al., 2013): object A approached object B (static), A made contact with B, and A stopped moving as object B continued along A’s same trajectory. Across trials for both types of events, objects varied in color (red, green, or blue), shape (rectangle or oval), and angle of trajectory (0 to 360 degrees, 45-degree increments). See Figure 2a for examples (statically depicted); dynamic versions may be viewed on our OSF repository (link on title page).

These visual stimuli had duration 1,250 ms and were displayed at a size of 640x400 px in the participant’s Web browser. Due to the nature of online studies, we cannot know the exact viewing distance, screen size, luminance (etc.) of these stimuli as they appeared to participants. However, any distortions introduced by a given participant’s viewing distance or monitor settings would have been equated across all stimuli and conditions.

Design and Procedure

The trial structure is schematized in Figure 2b. Participants were told to select the phrase that best matched the visual display. On each trial, participants viewed the same visual stimulus twice (either symmetrical or non-symmetrical), preceded by a fixation cross (350 ms). The stimulus disappeared and then a predicate pair appeared below it. Each predicate was preceded by infinitival “to” (e.g., “to chat” vs. “to tell”). Participants pressed the F or J key to make their choice, with mapping of predicate to keyboard key randomized on each trial (e.g., F for symmetrical, J for non-symmetrical). Note again that no predicate pair (besides the “collide”/“hit” item) was directly related semantically to the visual events. Trial order was pseudorandomized (24 trials total). Predicate pairs were randomly assigned to each stimulus without replacement.

Analysis

In accordance with our pre-registered analysis plan, individual trials with extremely fast RTs (less than 200 ms) were excluded. Participants were excluded for failure to contribute a complete dataset, reporting that they were not native English speakers, or responding with extremely fast RTs on more than 20% of trials. Five participants meeting these criteria were excluded, although none of the results reported here or in subsequent studies were dependent on the exclusion criteria.

Here and in subsequent experiments, we tested our predictions with mixed-effects logistic regression models on trial-level data, which allow for generalization of statistical inferences simultaneously across participants and items (Barr et al., 2013). The dependent variable was “symmetrical predicate choice”: choosing the symmetrical predicate (rather than non-symmetrical). The key independent variable was Visual Type (symmetrical vs. non-symmetrical, sum-coded as -0.5 and 0.5 , respectively). A main effect of Trial Number (centered) was included in the baseline model to account for general order effects, and its interaction with Visual Type was also tested in case the effect of interest changed over the course of the study.

Here and in all other mixed-effects model analyses reported in this paper, we tested for significance of variables by using likelihood ratio tests on the Chi-square values from nested model comparisons with the same

random effects structure. We started with the maximal random effects structure: correlated random intercepts and random slopes for Visual Type by participant and by item (predicate pair). If models did not converge, we simplified the random effects structure by first using uncorrelated intercepts and slopes, and followed that by dropping random intercepts and slopes until convergence, starting with those that accounted for the least variance.

For all experiments, we report effect sizes in terms of odds ratios (ORs), derived from β (the logit-transformed fixed effect coefficient). For the fixed effect of Visual Type, the OR represents the increased likelihood of a symmetrical predicate choice for one visual type vs. another. For example, an OR of 3.0 would mean that choosing the symmetrical predicate (vs. the non-symmetrical predicate) is three times more likely after observing a symmetrical collision event than an asymmetrical launch event.

The key prediction was a significant “matching” effect: for symmetrical choices to be higher for symmetrical vs. non-symmetrical visual stimuli (which would manifest as a main effect of Visual Type).

Results

Our prediction was confirmed: As can be seen in Figure 3a, when participants viewed visual events that were symmetrical, they more often selected an English predicate with a logically symmetrical meaning (e.g., “equal” or “meet”) than when they saw a visually non-symmetrical event, $\chi^2(1) = 19.90, p < .001$ (a main effect of Visual Type; $\beta = 1.13$ [95% CI 0.71 to 1.55], $z = 5.29, p < .001$, OR = 3.09 [95% CI 2.04 to 4.70]). In other words, symmetry across domains was “matched” in the minds of our participants. This was also evident non-parametrically: 50 of the 55 participants and 21 of the 24 predicate pairs went in the direction of this effect. Moreover, the effect was stable throughout the experiment: there was no significant interaction between Visual Type and Trial Number, $\chi^2(1) = 0.31, p = .576$. In fact, an exploratory analysis showed that the matching effect was already significant by the second trial ($\chi^2(1) = 7.80, p = .005$). Thus, symmetry was not only mapped across domains, but this mapping was also robust and relatively immediate.

Our hypothesis was that this correspondence is at an abstract level of symmetry. But as can be seen in Table 1, predicates varied from highly abstract (e.g., “to differ”) to highly concrete and observable (e.g., “to box”, in the sense of *to punch one another for sport*). And many events or situations that are described by symmetrical predicates just *look* visually symmetrical, like those described by the predicates *collide*, *tango*, and *box*. Were we just observing a match between what these events or situations typically look like, and our visual stimuli? To answer this, we conducted an exploratory analysis in which we extracted concreteness norms from Brysbaert et al., (2014), and examined whether these predicted the strength of the matching effect (i.e., the tendency to match symmetrical predicates to symmetrical visual stimuli, and non-symmetrical predicates to non-symmetrical visual stimuli). The rating scale ranged from 1 (least concrete) to 5 (most concrete). See Appendix A for these and other norms for the predicates we used.

Results of the concreteness analysis can be seen in Figure 3b. Remarkably, the matching effect held regardless of the concreteness of the predicate pair. This was confirmed using mixed-effects logistic regression, predicting “match” responses (symmetrical choices for symmetrical visual stimuli, and non-symmetrical choices for non-symmetrical visual stimuli). The average concreteness value (centered) of the predicate pair was not a predictor of the magnitude of the matching effect, $\chi^2(1) = 1.15, p = .284$. In fact, the effect of concreteness trended in the opposite direction: the higher the average concreteness of the predicate pair, the lower the matching effect ($\beta = -0.16$ [95% CI -0.45 to 0.13], $z = -1.09, p = .277$, OR = 0.85 [95% CI 0.64 to 1.14]). Thus, it was not merely concrete symmetrical predicates driving the matching effect; the correspondence appeared to be at an abstract level.

Figure 3. Results for Experiment 1a. (a) Matching results. Participants matched symmetry across linguistic predicates and visual events. Item means $\pm$ 95% CIs (within-item error bars). *** $p < .001$. (b) The matching effect was not explained by the concreteness of the predicates (norms from Brysbaert et al., 2014, with the scale ranging from 1 to 5): The effect was present whether predicates were concrete or abstract, with 21 of the 24 predicates going in the direction of the effect. (Only the symmetrical predicate for each pair is shown in this plot for simplicity of visualization.)

Experiment 1b: Symmetry Ratings for Predicates

Next, we collected ratings of the symmetry of predicates and related these to the matching effect of Experiment 1a. If the effect was driven by construal of symmetrical predicates as such, we should observe that symmetry ratings predict the degree to which participants matched a symmetrical predicate to a symmetrical visual stimulus. This would also allow us to confirm our decisions about the symmetry of predicates in independent data.

Participants

Thirty-two participants were recruited from MTurk. Seven of these participants were excluded for low catch-trial performance (see below). Demographic factors were the same as in Experiment 1a. Participants were paid \$1.00 for their participation, with an average study duration of seven minutes.

Stimuli, Design, and Procedure

Stimuli were the 48 predicates from Experiment 1a. In the instruction phase, participants were told that a word is symmetrical if its meaning is “mutual”, and they were given contrasting examples for explanation (“cousin” vs. “father”; “next to” vs. “on top of”). Participants were then presented with each predicate one at a time and were instructed to rate how symmetrical it was, from 1 to 6. For approximately half the participants, higher numbers on the scale indicated more symmetry (e.g., 6 corresponded to “very symmetrical”), and the opposite for the other half of participants (e.g., 1 corresponded to “not at all symmetrical”). The instruction phase examples (i.e., “cousin”, “father”, “next to”, “on top of”) were randomly interspersed among these test trials and were used as “catch” trials: if a participant gave less than 75% of these catch trials a rating at or one away from the appropriate symmetry rating (lowest or highest, depending on the particular item), they were excluded.

Analysis

Symmetry ratings for each predicate were z -scored within-participant and then averaged across participants. The key measure of interest was a “symmetry rating difference” within each predicate pair: the difference in average symmetry rating between a symmetrical predicate and its yoked non-symmetrical counterpart. To test how well symmetry ratings predicted the “matching effect”, we used the symmetry rating difference as a predictor in mixed-effects logistic regression on trial-level data from Experiment 1a. Specifically, the dependent variable was “match choice”: choosing the predicate that matched the symmetry of the visual stimulus (“symmetrical choice” for visual symmetry trials, and “non-symmetrical choice” for visually non-symmetrical trials). Independent variables were Visual Type (symmetrical vs. non-symmetrical, sum-coded as -0.5 and 0.5 , respectively) and Symmetry Rating Difference (centered). We predicted a significant main effect of Symmetry Rating Difference.

Figure 4. Results for Experiment 1b. (a) Mean symmetry ratings for individual predicates on the y-axis. Symmetrical predicates (blue bars) were overwhelmingly rated as more symmetrical than the non-symmetrical predicates (orange bars), with only a few exceptions near the middle of the scale. (b) The symmetry rating differences (the differences in average rating between a symmetrical predicate and its non-symmetrical foil) appear on the x-axis in each panel. These rating differences predicted the Experiment 1a “matching” effect (on the y-axis), more so for symmetrical visual events ($p < .001$, the left panel in blue) than non-symmetrical events ($p = .578$, the right panel in orange), with a marginal interaction between the two ($p = .106$). (Only the symmetrical predicate for each pair is shown in this plot for simplicity of visualization.)

Results

First, we confirmed that the symmetry ratings collected in this study were highly reliable. The split-half reliability between even- and odd-numbered participants was near ceiling (Spearman’s $\rho = 0.94$). And as can be seen in Figure 4a, our own predicate selections were validated by these ratings: the vast majority of symmetrical

predicates were rated as high on symmetry and non-symmetricals as low. There were a few notable exceptions: the non-symmetrical predicates *befriend*, *contact*, *greet*, and *consent* were at or above average on symmetry, while the symmetrical predicates *box* and *differ* were at or below average. Crucially, however, within predicate pairs, the symmetrical predicates were conceived of as more symmetrical than their non-symmetrical counterparts: All predicate pairs had a positive ζ -scored Symmetry Rating Difference ($M=1.23$, $SD=0.56$), ranging from “correspond”/“contact” (0.12) up to “be identical”/“be inferior” (2.42). Mean ratings for each predicate can be found in Appendix A.

As predicted, the difference in symmetry ratings correlated with the matching effect: a mixed-effects logistic regression that included the factor Symmetry Rating Difference was a better fit than one with only a main effect of Visual Type, $\chi^2(1) = 11.68$, $p < .001$. Surprisingly, as can be seen in Figure 4b, this relationship was stronger for symmetrical collisions than asymmetrical launches: we observed a marginal improvement for a model with the interaction of Symmetry Rating Difference and Visual Type over one without this interaction, $\chi^2(1) = 2.61$, $p = .106$. In separate analyses on each visual type, we found that for collision trials, a model with the factor Symmetry Rating Difference was a significantly better fit than one without ($\chi^2(1) = 13.14$, $p < .001$; Rating Difference $\beta = 0.81$ [95% CI 0.42 to 1.20], $\zeta = 4.06$, $p < .001$, OR = 2.24 [95% CI 1.52 to 3.31]). This was not the case for launch trials ($\chi^2(1) = 0.31$, $p = .578$; Rating Difference $\beta = 0.16$ [95% CI -0.40 to 0.72], $\zeta = 0.56$, $p = 0.575$, OR = 1.17 [95% CI 0.67 to 2.05]).

Our interpretation of this difference between visual types is the following (which we note is necessarily speculative). We suspect that when observers view visually symmetrical stimuli such as those in Experiment 1a, these stimuli inextricably evoke a notion of symmetry. This makes symmetry available and salient as a dominant factor for matching, beyond the other rich semantic information these predicates also convey. By contrast, non-symmetrical visual stimuli do not strongly evoke symmetry (or even an absence of symmetry), leaving other semantic properties available for matching.⁴

Experiment 2: One vs. Two Static Objects

Experiment 1a demonstrated that a correspondence exists between logically symmetrical predicates and visually symmetrical stimuli, and Experiment 1b showed via symmetry ratings of the predicates that this correspondence was driven by symmetrical construal per se (as the strength of the matching effect was predicted by the ratings). These initial findings raise intriguing new questions about the nature of this link, which we probed in the current study.

First, we asked whether dynamic stimuli such as the visual events of Experiment 1a are necessary to elicit such a correspondence. There is reason to think not. After all, symmetrical predicates can be used to refer not just to events (e.g., “to tango”, “to collide”), but also to states (e.g., “to match”, “to differ”). Likewise, the visual symmetry of static figures is a core property extracted automatically in visual processing (Wagemans, 1997).

Second, we asked whether construal of the visual stimulus as a binary relation between entities (i.e., “relational symmetry”) makes this correspondence more salient. There are reasons to think it might. Recall that

⁴ We also note that our symmetrical predicates tended to be more positively valenced or “pleasant” than non-symmetrical predicates (e.g., compare *identical* vs. *inferior*, or *intersect* vs. *interfere*). However, the effect of symmetry on the matching effect could not be fully explained by this valence difference: Even though there was a marginally significant difference between our symmetrical and non-symmetrical predicates on affective valence norms (Warriner et al., 2013), $t(23) = 1.95$, $p = .062$ (see Appendix A for values on these norms), a mixed-effects model with the main effect of Symmetry Rating Difference predicted the matching effect over and above a model with just the Valence Difference (centered) on its own, $\chi^2(1) = 7.42$, $p = .006$.

in language, symmetry is a property of binary relations: i.e., relations holding between two *discrete entities*. It turns out that certain visual processes such as attention also operate over discrete entities — visual objects — and not just over features (e.g., edges, parts) or spatial regions (for a review, see Scholl, 2001). Remarkably, this even extends to visual processing of certain kinds of relations. For example, chasing is better detected when it occurs between discrete objects rather than object *parts* (van Buren et al., 2017) — even though such parts can be *cognitively* construed as discrete entities. Indeed, the visual events of Experiment 1a featured two discrete objects as event participants and elicited a strong matching effect. Nevertheless, the extent to which observers analyze visual relations between static objects as symmetrical is unknown (cf. Baylis & Driver, 2001), although recent work suggests that under the right conditions, young children may process between-object symmetry as explicitly relational, rather than as a purely low-level feature (Shao & Gentner, 2019).

Here we asked whether the linguistic-visual correspondence for symmetry extends to static objects that convey symmetry, and we tested whether such a correspondence is stronger for two discrete visual objects, where perceptual accessibility of the relation's binary nature is more salient.

Participants

One hundred participants were recruited from the online platform Prolific. We chose this sample size based on pilot data, which revealed that we would need a larger sample size than that of Experiment 1a to find the predicted interaction effect between visual symmetry and number of visual objects. Demographic factors and exclusion criteria were the same as in Experiment 1a (one participant was excluded). Participants were paid \$0.84 for their participation, with an average study duration of six minutes.

Stimuli

Linguistic stimuli were the predicate pairs from Experiments 1a and 1b. Visual stimuli were simple objects constructed using methods similar to Baylis and Driver (2001). Objects were eight rectangular blocks (pseudorandom width) connected adjacently, top to bottom. Each object was surrounded by a black border and was placed with a slight shadow on a textured background; together these cues bias interpretation of the objects as being the perceptual figure rather than ground. Objects were filled in red, green, or blue. Examples appear in Figure 5a. Twenty-four different exemplars of such objects were created.

Two factors were crossed in this experiment. First, visual stimuli were either symmetrical or non-symmetrical (the Visual Type factor). Second, visual stimuli were either one or two objects (the Object Number factor). For single symmetrical objects, rectangular blocks were mirror-reflected about the vertical axis. Two-object versions of each single object were created by making two differently-colored objects face each other with the same inward-facing contour as the contour of the single symmetrical object. Non-symmetrical versions of each stimulus were made by offsetting 75% of the blocks horizontally by a pseudorandom number of pixels.

Design, Procedure, and Analysis

Design, procedure, and analysis were similar to Experiment 1a, but with one additional factor: Object Number (One or Two). There were two blocks of 12 trials each: a One-Object block and a Two-Object block (with block order counterbalanced across participants). Within-block, trials were evenly divided between symmetrical and non-symmetrical visual stimuli (trial order randomized). Block Order (sum-coded) was included as an additional factor in model comparison, in case main effects or interactions differed by block order. Images were displayed at 536x402 px for 2,784 ms and preceded by a 350 ms fixation cross.

Once again, the key prediction was a significant “matching” effect, i.e., for symmetrical choices to be higher for visually symmetrical vs. non-symmetrical stimuli (which would manifest as a main effect of Visual Type). We also predicted an interaction between Visual Type and Object Number, whereby the “matching” effect for the Two-Object condition would be greater than the One-Object condition.

Results

Data appear in Figure 5b. First, as in Experiment 1a, participants matched symmetrical predicates to symmetrical visual stimuli more often than to non-symmetrical visual stimuli, $\chi^2(1) = 13.24, p < .001$ (a main effect of Visual Type). Confirming our second prediction, the matching effect was stronger for Two-Object than One-Object trials. This manifested as a significant interaction of Visual Type and Object Number, $\chi^2(1) = 4.76, p = .029$, compared to a model with only main effects of these factors (Visual Type $\beta = 0.82$ [95% CI 0.44 to 1.21], $z = 4.18, p < .001$, OR = 2.28 [95% CI 1.55 to 3.35]; Object Number $\beta = 0.34$ [95% CI 0.17 to 0.51], $z = 3.89, p < .001$, OR = 1.41 [95% CI 1.19 to 1.67]; β for their interaction = 0.38 [95% CI 0.04 to 0.73], $z = 2.18, p = .029$, OR = 1.47 [95% CI 1.04 to 2.07]).

As in Experiment 1a, there were no significant order effects: the Two-Object advantage for the matching effect held whether participants viewed the One-Object or Two-Object block first (no significant interaction of Block Order, Visual Type, and Object Number, $\chi^2(3) = 0.84, p = .839$). Additionally, an exploratory analysis showed that the main effect of Visual Type was already significant by the fourth trial ($\chi^2(1) = 11.43, p < .001$), and its interaction with Object Number approached significance by the halfway point of the experiment ($\chi^2(1) = 3.70, p = .054$).

Finally, as for the visual event stimuli of Experiment 1a, we explored whether symmetry ratings of predicates predicted the strength of the matching effect in the current experiment. A mixed-effects logistic regression that included the factor Symmetry Rating Difference was a better fit than one with only main effects and interactions of Visual Type and Object Number, $\chi^2(1) = 6.75, p = .009$. This relationship was stronger for symmetrical visual stimuli than non-symmetrical visual stimuli: we observed a marginal improvement for a model with the interaction of Symmetry Rating Difference and Visual Type over one without this interaction, $\chi^2(1) = 2.69, p = .101$. In separate analyses on each visual type, we found that for visually symmetrical trials, a model with the factor Symmetry Rating Difference was a significantly better fit than one without ($\chi^2(1) = 9.27, p = .002$; Rating Difference $\beta = 0.63$ [95% CI 0.26 to 1.00], $z = 3.36, p < .001$, OR = 1.88 [95% CI 1.30 to 2.72]). This was not the case for visually non-symmetrical trials ($\chi^2(1) = 0.88, p = .349$; Rating Difference $\beta = 0.20$ [95% CI -0.22 to 0.62], $z = 0.94, p = 0.346$, OR = 1.22 [95% CI 0.80 to 1.87]). We also explored whether the concreteness of the predicates would predict the strength of the matching effect. As in Experiment 1a, the average concreteness value (centered) of the predicate pair was not a predictor of the magnitude of the matching effect, $\chi^2(1) = 1.10, p = .293$, and in fact trended in the opposite direction: the higher the average concreteness of the predicate pair, the lower the matching effect ($\beta = -0.16$ [95% CI -0.45 to 0.13], $z = -1.06, p = .289$, OR = 0.85 [95% CI 0.64 to 1.14]). Thus, as in Experiment 1a, it was not merely concrete symmetrical predicates driving the matching effect here; again the correspondence appeared to be at an abstract level.

To summarize, the results from this experiment suggest two primary conclusions: first, the linguistic-visual correspondence for symmetry extends to static visual objects, and second, by default the mind analyzes two symmetrically configured visual objects in terms of a binary relation, which can then more readily be associated with symmetrical predicates.

Figure 5. Experiment 2 stimuli and results. (a) Stimuli were static figures (symmetrical or non-symmetrical) generated by placing stacked rectangles of different widths on a textured background. Crucially, the stimulus set included stimuli in which the same symmetrical (or non-symmetrical) contours were present for the one-object and two-object conditions. (b) Participants matched symmetry across predicates and static objects, more so for two objects than one. Item means $\pm 95\%$ CIs (within-item error bars). *** $p < .001$; ** $p < .01$; * $p < .05$.

Experiment 3: Linguistic Signs

Our results demonstrate a deep connection between representations of symmetry across language and vision. These observations raise the intriguing possibility that we might find traces of this relationship within natural language itself, in languages that happen to be able to express abstract relations spatially, like sign languages.

Remarkably, there is already some evidence for an intuitive linguistic-visual link for symmetrical relations. Recent work has demonstrated that speakers of Nicaraguan Sign Language (NSL), an emerging language among deaf individuals, spontaneously express logically symmetrical concepts such as “high-five” spatially, by using signs with mirror-reflective symmetry (Gleitman et al., 2019).

Does such a linguistic-visual link also hold for non-signing observers, and if so, what is the nature of this link? In a final study, we exploited the tendency of sign languages to express relations spatially to look for further evidence of symmetry as a unit of meaning intuitively linked to symmetrical visual forms. We were inspired by Strickland et al. (2015), who investigated the sensitivity of observers to visual reflexes of *telicity*, which is a grammatical property of verb phrases that indicate an inherent endpoint or culmination for the events to which they refer (e.g., contrast *John thought for an hour* with the infelicitous *John decided for an hour*). Strickland and colleagues demonstrated in a cross-modal matching paradigm that naive observers with no knowledge of a sign language nevertheless appear to have access to a possibly universal mapping bias between visual cues to telicity (gestural boundaries, present across diverse sign languages) and grammatical telicity (in their own language, English), such that they tended to select telic English verbs (e.g., *decide*, *enter*) for telic signs and atelic verbs (e.g., *think*, *run*) for atelic signs, when forced to choose.

We took the same approach here. Observers who were unfamiliar with sign languages were tasked with guessing the English translations of signs from American Sign Language (ASL). Crucially, the ASL signs were never actually a direct translation of the English phrases presented on a trial. We asked whether non-signers would select symmetrical English predicates for visually symmetrical signs and non-symmetrical English predicates for non-symmetrical signs. (See below for how signs were defined as visually symmetrical.) If so, this would provide evidence that such a mapping bias exists between conceptual symmetry and visual form in natural language itself.⁵ It would also demonstrate that the linguistic-visual correspondence for symmetry can be observed despite the perceptual richness and variety of visual signs — going beyond the controlled but constrained visual stimuli used in our previous studies.

Participants

Sixty participants were recruited from the online platform Prolific. We chose this sample size to match that used in Experiment 1a. Demographic factors and exclusion criteria were the same as in Experiments 1a and 2, with the addition of criteria related to knowledge about sign languages (both ASL and others). In particular, we excluded any participant who reported having even basic knowledge of a sign language. All exclusion criteria were pre-registered, as in our previous matching studies. Participants were paid \$1.00 for their participation, with an average study duration of 7 minutes.

Two participants were excluded for too many fast RTs (> 20% of trials with RTs < 200 ms), and 16 were excluded for reporting having any knowledge of a sign language, leaving 42 participants. (Results were qualitatively the same with or without these excluded participants.)

⁵ It is important to note that visually symmetrical signs exist in ASL which do not convey a relationally symmetrical meaning (e.g., the signs for *expand* or *consent*). However, among ASL verbs, there does seem to be a form-meaning correspondence, even if probabilistic. For example, an inspection of ASL dictionaries revealed that the ASL translations of our symmetrical English predicates were more likely to have a visually symmetrical sign than translations of non-symmetrical predicates: 13 symmetrical predicates had ASL signs that were mostly or fully visually symmetrical (and all but one were two-handed), while only 6 non-symmetrical predicates did (with an additional five being one-handed signs). Future work would need to establish this form-meaning correspondence empirically, with a systematic examination of the meanings of these signs based on ASL signer judgments of the linguistic-semantic criteria of symmetry, rather than relying on dictionary translations alone.

Stimuli

English linguistic stimuli were the predicate pairs from the previous studies. To select ASL signs, we wanted to choose 24 signs (12 symmetrical, 12 non-symmetrical) whose visual symmetry matched the symmetry of their English translations. When possible, we attempted to include ASL signs that corresponded to the English predicates we used; crucially, however, the ASL sign on a given trial was never actually a direct translation of the two choices for English predicates. For example, the sign for the English predicate *collaborate* is visually symmetrical, so we included this in our ASL stimulus set; however, this sign was never presented with the English predicate pair “to collaborate” vs. “to contribute” as options; instead, the ASL sign for *collaborate* might be randomly paired with the predicate pair “to date” vs. “to befriend”.

To find candidate signs, we first looked for ASL sign translations of our English predicates in the online ASL databases www.handspeak.com and www.signingsavvy.com. Then in consultation with a hearing ASL signer (a “coda”, or child of deaf adult), we reviewed these signs to confirm that the translations were correct and to characterize the visual properties of these signs, as follows. We categorized ASL signs as visually symmetrical if they used two hands with shapes and movements that were reflectively symmetrical in the vertical plane. We categorized ASL signs as visually non-symmetrical if they involved the two hands but did not demonstrate reflective symmetry of hand shapes or movements. Fourteen ASL signs (7 symmetrical and 7 non-symmetrical) matching our English predicates met these criteria unambiguously. (That only some of the ASL signs did so is not surprising, given our strict definitions of visual symmetry; many more met some but not all of these criteria, e.g., they had symmetrical movements but not hand shapes.) To select another 10 ASL signs, we searched for ASL translations of common English symmetrical and non-symmetrical predicates meeting the linguistic criteria for symmetry (or non-symmetry) outlined in the Stimuli section of Experiment 1a, until we found a candidate set of 30 ASL signs, in total.

After building this candidate list, we produced videos of these signs for our study. Another bilingual English-ASL speaker (also a coda) naive to the purposes of the study was filmed producing these signs in presentation form. The speaker wore a black shirt, maintained a neutral facial expression, and produced the signs in front of an off-white curtain. These videos were edited (with black buffer on either side), and a final 24 signs were selected from this set (with the duration of videos between 3 and 4 seconds long). The visual symmetry (or non-symmetry) of this final set of signs was confirmed in a separate norming study.⁶ The list of English translations of the ASL signs used appears in Table 2, and example still images from the signs appear in Figure 6a. Full videos are available on our OSF repository (link on title page).

⁶ Thirty-eight undergraduate participants from the University of Pennsylvania rated the symmetry of these ASL signs in an online study for course credit. Participants were instructed to rate how visually symmetrical each sign was, one at a time, from 1 to 6. (For instruction, they first viewed an example of a symmetrical and an asymmetrical butterfly). The 12 symmetrical signs had a mean symmetry rating of 5.12 ($SD = 1.17$, range 3.63–5.92), while the 12 non-symmetrical signs had a mean symmetry rating of 1.87 ($SD = 1.19$, range 1.21–3.26).

Table 2: English translations of the ASL signs (Symmetrical and Non-Symmetrical) used in Experiment 3. Those marked with “x” are ASL signs that directly translate to one of the English predicates in the list of predicate pairs in Table 1 (used in all of our experiments).

Symmetrical		Non-Symmetrical	
ASL Sign	In Predicate List	ASL Sign	In Predicate List
1. argue		13. assist	
2. box	x	14. crash	
3. collaborate	x	15. drown	
4. debate	x	16. expand	x
5. diverge		17. hit	x
6. divorce		18. be inferior	x
7. embrace		19. intervene	x
8. equal	x	20. lead	x
9. match	x	21. reject	x
10. negotiate	x	22. seduce	
11. quarrel		23. touch	
12. unite	x	24. be typical	x

Design, Procedure, and Analysis

Design, procedure, and analysis were similar to Experiment 1a, except that the task and instructions differed: instead of participants picking which stimulus was the “best match”, participants were asked to choose which English phrase was the one that the person was signing. In other words, they were told to guess the meaning of the sign. Crucially, the ASL signs were never a direct translation of the English phrases presented on a trial. For example, the ASL sign for *boxing* (i.e., to punch one other for sport) would never be presented with the English predicate pair “to box” vs. “to punch”. Thus, we could be confident that any other visually transparent aspect of the sign’s meaning apart from symmetry (e.g., fists for boxing) would not lead to a preference for the symmetrical English predicate over the non-symmetrical one. Predicate pairs were randomly paired with ASL signs while heeding this constraint.

Within-block, trials were evenly divided between Symmetrical and Non-Symmetrical visual stimuli (with trial order randomized). Videos were displayed twice (at 576x360 px), preceded by a white 300 ms fixation cross on a black background each time. Then the English predicate pair choices appeared.

Once again, the key prediction was a significant “matching” effect, i.e., for symmetrical choices to be higher for visually symmetrical ASL signs as compared to visually non-symmetrical ASL signs (which would manifest as a main effect of Visual Type).

Results

Data appear in Figure 6b. As in Experiments 1a and 2, participants matched symmetrical predicates to symmetrical visual stimuli more often than to non-symmetrical visual stimuli, $\chi^2(1) = 11.01, p < .001$ (a main effect of Visual Type; $\beta = 0.60$ [95% CI 0.28 to 0.91], $\zeta = 3.72, p < .001$, OR = 1.82 [95% CI 1.33 to 2.50]). This was also evident non-parametrically: 30 of 42 participants and 18 of 24 predicate pairs went in the direction of this effect. Moreover, the effect was stable throughout the experiment: there was no significant interaction between Visual Type and Trial Number, $\chi^2(1) = 1.58, p = .209$. Indeed, an exploratory analysis showed that the effect of Visual Type was already significant by the tenth trial ($\chi^2(1) = 6.23, p = .013$).

Finally, as in the previous experiments, we explored whether symmetry ratings of predicates predicted the strength of the matching effect in the current experiment. A mixed-effects logistic regression with the factor Symmetry Rating Difference was not a significantly better fit than one without ($\chi^2(1) = 2.21, p = .137$), although it did trend in the expected direction (Rating Difference $\beta = 0.20$ [95% CI -0.06 to 0.47], $\zeta = 1.52, p = .129$, OR = 1.23 [95% CI 0.94 to 1.60]). We also explored whether the concreteness of the predicates would predict the strength of the matching effect in this study. As in previous experiments, the average concreteness value (centered) of the predicate pair was not a significant predictor of the magnitude of the matching effect, $\chi^2(1) = 0.19, p = .664$.

We can conclude several things from these findings. First, the tendency to match symmetry across the visual and linguistic domains extends beyond simplified events and objects, to rich and naturalistic stimuli. Furthermore, this tendency extends beyond an artificial matching paradigm, to one in which observers were asked to infer the meaning of words in a sign language in which relational symmetry is often expressed spatially. Thus, the mapping between the symmetry of predicates and of the visual form of signs is intuitive — just as Strickland et al. (2015) found for telicity and visual boundedness.

Figure 6. Experiment 3 stimuli and results. (a) Visual stimuli were signs from American Sign Language (ASL), symmetrical and non-symmetrical. These stimuli were dynamic; example still frames appear here, with previous hand positions transparently overlaid and with white arrows showing the direction of hand movements for illustration (arrows were not present in the videos shown to participants). (b) Participants matched symmetry across English predicates and ASL signs. Item means $\pm 95\%$ CIs (within-item error bars). *** $p < .001$; ** $p < .01$; * $p < .05$.

General Discussion

The four experiments reported here show that mental representations of symmetry correspond across cognitive systems, and our results reveal several aspects of the nature of this correspondence. First, individuals have access to the symmetry of concepts tokened by English predicates. Second, individuals are sensitive to visual cues to symmetry. Third, the commonality in the type of symmetry accessed by both vision and language is intuitive, such that individuals in our studies associated symmetrical items across these two domains. This was true even though the predicates (e.g., *equal*, *differ*) did not in general relate transparently to the visual stimuli — nor indeed to visual concepts at all.

We also identified the generality of this mapping, and some of its constraints. It held not only for dynamic events (Experiment 1a), but also static figures (Experiment 2) and even rich visual signs (Experiment 3). Furthermore, it generalized across tasks: whether forced to choose a best “match” for a simplified visual stimulus (Experiments 1a and 2) or to guess the meaning of an ASL sign rich in visual information (Experiment 3), participants exhibited similar matching effects. And finally, making the visual stimuli’s binary nature less apparent (as in the single-object condition of Experiment 2) weakened participants’ ability to recognize this correspondence. Apparently, symmetry across linguistic and visual systems is most obvious when it is “relational”, i.e., when it holds for a binary relation, in both cases. This finding is consistent with previous work showing qualitative differences in processing within- and between-object visual symmetry (Baylis & Driver, 2001; Shao & Gentner, 2019).

Where Word and World Meet: Approaching Concepts “From Both Sides”

Mathematicians define symmetry as invariance under transformation (Weyl, 1952), which is a formal characterization that may be applied to symmetry in both the visuospatial and linguistic domains. However, our results go beyond purely formal description, by demonstrating that the mind itself actually *treats* symmetry in the visual and linguistic domains as similar, in ways sufficient to engender a correspondence cross-modally. Given that stimuli in these domains share no surface features, we suggest that both vision and language point to a common, amodal form of symmetry in the conceptual system, where representations of symmetry from both domains may be directly compared (though for alternative possibilities, see subsection below entitled “Possible Mechanisms of Cross-Modal Matching”). How is this “central” form of symmetry accessed “from both sides”?

On the language side, we propose that linguistic predicates token non-linguistic relational concepts that have the property of symmetry. Importantly, our findings demonstrate that the mind has access to this symmetry even when such predicates are considered in isolation (e.g., *to marry*), without requiring that they be embedded in structured linguistic utterances (e.g., *Mark and Bill marry*). That notions of symmetry are elicited by isolated lexical items suggests that symmetry might even be a lexical-semantic property, whereby some predicates but not others are “marked” linguistically as belonging to the symmetrical class, and so enjoy the relevant interpretive privileges (e.g., *John and Bill met* is interpreted reciprocally, and *John met* — though syntactically permitted — is semantically anomalous) (Gleitman et al., 1996). Our results thus provide even stronger support for Gleitman et al.’s (1996) original claim, contra Tversky (1977), that the asymmetry sometimes attributed to linguistic terms for symmetrical concepts such as *similar* is introduced by other aspects of the linguistic system; namely, the asymmetrical syntactic structures in which they are embedded (e.g., *Mark is similar to Bill*). Nevertheless, we recognize that single lexical items, including symmetrical ones, interact with syntax in rich and systematic ways in utterance interpretation (Gleitman and Landau 2015; see Williams, 2021, for an overview of how different types of events may be characterized semantically in structured linguistic utterances).

On the vision side, we propose that high-level perception generates representations of symmetry that are abstract and relational, such that they may easily interface with representations of symmetry in the conceptual system (Quilty-Dunn, 2020). Two aspects of our data support this proposal. First, the perceptual symmetry tapped by our task was abstract. Across experiments, we did not observe a relationship between concreteness and the matching effect (i.e., the matching effect was no stronger — or maybe even a bit weaker — for predicates such as *collide* or *tango* that describe visually symmetrical events than for abstract items like *differ* or *negotiate*; see Figure 3b). Second, in Experiment 2 (static shapes), the matching effect was stronger for two symmetrically configured objects than for one object with a symmetrical contour. Given that discrete objects rather than lower-level features are the units of higher-level vision (Scholl, 2001; van Buren et al., 2017), we can conclude that a binary relation with the property of symmetry is generated by high-level vision and furnished to cognition.

Our work adds to a growing literature showing that perceptual processing goes beyond low-level properties such as the colors, shapes, locations, or motions of objects. Instead, perception in some cases appears to generate representations of relations at an abstract, structured level (for a review, see Hafri & Firestone, 2021). This project also joins an emerging literature exploring the interface between language and visual cognition, and how they share surprisingly sophisticated content and representational principles (Cavanagh, 2021; Hafri et al., 2018, 2020, 2022; Strickland, 2017). Of course, one caveat to our claims about the visual system’s precise role in the linguistic-visual mapping is that a stimulus being visual does not on its own entail that all of its high-level properties (such as symmetry) are perceived via rapid, automatic visual processes. Instead, one might infer the existence of a property by *reasoning* over more basic visual features, such as contours or motion trajectories. Although our work is consistent with the possibility that perceptual processing of symmetrical relations is automatic, future work may confirm and enrich this conclusion by designing cross-modal matching tasks that implicate automatic perceptual processing (e.g., occurring rapidly and without effort or special instruction; Fodor, 1983; for more recent perspectives, see Scholl & Gao, 2013, and Hafri & Firestone, 2021).

Possible Mechanisms of Cross-Modal Matching

We believe that the comparison between representations elicited by the visual and linguistic stimuli in our task requires a common representational format, and we assume that this format is abstract and amodal. Nevertheless, some readers may question this assumption. Perhaps instead the representations of symmetry being compared in our study are rooted directly in sensorimotor systems — including the representations elicited by linguistic stimuli. Although we do not address the wider debate between abstract and sensorimotor or “embodied” accounts of concepts here (for examples, see Glenberg 2015; Mahon, 2015a, 2015b), there is one obvious prediction from such a view: If the mapping were based only on concrete sensorimotor representations (elicited by both visual and linguistic stimuli), then one might expect the matching effect to be especially strong for predicates that describe visually symmetrical events such as *collide*, *tango*, or *chat*. By contrast, across our experiments we found no significant effect of concreteness, and we even found in some cases that the effect trended in the opposite direction. These analyses suggest that the representations of symmetry being compared in our studies are not purely sensorimotor in nature.

Another alternative account is that observers in our task formed analogies between the visual and linguistic stimuli directly, in which case a common (amodal) intermediate representation need not be posited at all. Of course, this process would still require distinct symmetry representations to exist *a priori* in both cognitive systems; these representations would then be placed into correspondence via processes of structural alignment (Gattis, 2004; Goldwater & Gentner, 2015; Goldwater et al., 2011; Markman & Gentner, 1993). Although it is

hard to imagine how to empirically distinguish this account from our own, we think that some aspects of our data argue against an analogical reasoning process, at least of an explicit and deliberate sort. In particular, it has been suggested that the structural alignment process can be computationally intensive, such that more time may be needed to represent and compare stimuli with relational properties than mere surface similarities (Markman & Gentner, 2000; Ratcliff & McKoon, 1989). By contrast, the cross-domain correspondence in our experiments emerged extremely quickly, both across trials and within individual trials, as we describe below.

Specifically, the effect emerged early-on in the experiments, often after just a few trials. Crucially, it did so despite that so much else besides symmetry differed between symmetrical and non-symmetrical English predicates (and likewise between the visually rich symmetrical and non-symmetrical ASL signs of Experiment 3). Thus, our results are evidence that symmetry was an especially prominent dimension of similarity, overshadowing the many other possible ways the visual and linguistic stimuli may in principle have been compared. Second, additional exploratory analyses suggest that the cross-modal matching effect was no stronger — or even a bit weaker — on trials where participants took longer to respond (indexed by response times normalized within-participant): In Experiment 1 (visual events), longer RTs predicted a weaker matching effect (normalized RT $\beta = -0.13$, $\chi = -1.91$, $p = .056$); a similar (but weaker) trend was observed in Experiment 3 (ASL signs; $\beta = -0.09$, $\chi = -1.34$, $p = .179$); and no relationship was observed in Experiment 2 (static objects; $\beta = -0.02$, $\chi = -0.33$, $p = .739$). Thus, to the extent that relatively longer latencies are indicative of more explicit reasoning or reflection, it appears that such reflection has no effect on, or may actually *diminish*, the influence of symmetry on matching responses — perhaps because participants begin to consider less obvious stimulus connections to make their judgments (e.g., the color or texture of the visual stimuli, or the contextual settings where certain events take place). Of course, future work may implement response deadlines to causally probe the latencies at which cross-domain matching effects for symmetry are greatest (e.g., Markman & Gentner, 2000; Ratcliff & McKoon, 1989).

Finally, Gleitman et al.'s (2019) recent empirical findings that the formal category of symmetry is present and expressed spatially in an emerging sign language (NSL) provide compelling evidence that the link between representations of symmetry in the linguistic and visual systems has its basis in an amodal conceptual representation of symmetry, rather than in purely sensorimotor representations or explicit analogical processes. To understand why, consider the “projectibility problem” posed by symmetrical predicates (Goodman, 1955): the learner must infer that symmetry holds for all possible pairs (x,y) in a given relation despite exposure to only a finite number of instances (e.g., $R(x_1,y_1)$ and $R(y_1,x_1)$, $R(x_2,y_2)$ and $R(y_2,x_2)$, etc.). Crucially, because the first NSL signers did not have access to a language model (i.e., someone from whom to learn an established natural language), an external linguistic-visual mapping was not yet available as a cue to facilitate the requisite inductive inferences for symmetrical predicates. Consequently, it must be the case that conceptually *construing* certain events or situations as symmetrical comes first (e.g., those described by predicates such as *equal* or *differ*), which then licenses their symmetrical spatial expression.

In sum, rather than the symmetry of linguistic predicates being *constituted* by sensorimotor patterns of symmetry, or being related to visual stimuli by a process of analogy, our results (and those of Gleitman et al., 2019) point to a rich *interaction* between abstract conceptual systems and sensorimotor systems (Mahon, 2015b), such that perceptual systems make available surprisingly abstract and structured representations to conceptual systems — here, for symmetry itself.

Implications for Learning Symmetrical Predicates

Our participants were native English-speaking adults, who have had a lifetime of hearing and using symmetrical predicates, leaving open the question of how children acquire these terms. As Gleitman et al. (2019) showed

with the spontaneous emergence of symmetry in NSL, the abstract notion of symmetry is *available* to the mind, even before knowledge of individual symmetrical lexical items develops, and this symmetry is formal in nature, generalizing to all instances of a given type (e.g., all events referred to by the term for *high-fiving*). In other words, symmetry for such predicates is projectible to all future possible instances (Goodman, 1955). Thus, this is at its core a mapping problem: discovering which phonological forms encode which (abstract) meanings.

How does this knowledge develop? For example, take two words “hit” and “collide”, that are very similar in meaning, apart from symmetry; how is a child to infer that just one of them, “collide”, points to a symmetrical concept? Our results, if they generalize to young children, suggest one possible learning story: There might exist perceptual “gems” for observable symmetrical situations (e.g., colliding, shaking hands, mutual hugging) that could enable the child to acquire some of the more concrete symmetrical words, if they happen to be uttered in such contexts. For example, if the child hears “Look, colliding!” in the presence of a canonical collision event, they might infer that *collide* is symmetrical. In other words, the mapping from perceptual to conceptual symmetry would constrain the hypothesis space of possible meanings for such a predicate.

Nevertheless, for more abstract words such as *marry* or *similar*, observation alone, however sophisticated, cannot be sufficient; thus, the symmetry of such abstract predicates must be acquired via other means. For example, via a learning procedure known as syntactic bootstrapping (Gleitman, 1990; Landau & Gleitman, 1985; see also Fisher et al., 2019), children could use knowledge of the syntactic restrictions or privileges of known symmetrical items to acquire unfamiliar items. Recall the two-part “litmus test” discussed in the Introduction: the singular subject intransitive (e.g., *Mark marries*) is semantically anomalous, and the intransitive has a roughly reciprocal interpretation (*Mark and Bill marry* is interpreted as Mark and Bill marrying *each other*). Perhaps children are sensitive to these properties and use them for learning which terms have symmetrical interpretations.

However, the ultimate solution will certainly prove complex, as there is no one syntactic structure in English unique to logical symmetry (see examples (1) to (5)). For example, the listener might notice when a speaker uses different syntactic alternations (1) to (4) in a single discourse to refer to the same event (e.g., “Look, Mark and Bill are hugging! Mark hugged Bill.”) (Yuan & Fisher, 2009). Or perhaps the listener might notice a suspicious absence of a singular subject in the intransitive, as in (5), for certain predicates. We suggest that investigating the observational and linguistic cues to symmetry (and how they might be used in concert) is a fruitful avenue for future research (e.g., Miller, 1998).

Beyond Language and Vision: Symmetry in Mathematical Reasoning

Intriguingly, the relationship between visual symmetry and abstract thought may extend to other forms of transformation invariance beyond the concepts tokened by predicates in natural language. A recent paper explored the relationship between the visual symmetry of mathematical symbols and the related mathematical concept of commutativity (i.e., where order of operations is interchangeable, as for addition and multiplication) (Wege et al., 2020). The authors found that participants have a strong tendency to associate vertically symmetrical symbols (e.g., ▲) with commutativity, as compared to when these same symbols were rotated 90 degrees (e.g., ►), which is known to diminish the salience of visual symmetry (Wagemans, 1997). Strikingly, they found that this association predicted participants’ performance on a mathematical test using these arbitrary symbols. Thus, the intuitive link between visual symmetry and abstract properties of symmetry may be quite general in nature, extending to the concepts used in abstract mathematical reasoning.

Relatedly, it is well known that many young children struggle to obtain an adequate understanding of mathematical equivalence in the early school years, leading to later difficulties learning algebra (e.g., Johannes

& Davenport, 2017; McNeil, 2014; National Council of Teachers of Mathematics, 2000). These difficulties have been connected to the pedagogical tendency to put the operators to the left of the equal sign (e.g., “ $2+2=4$ ” rather than also “ $4=2+2$ ”), perhaps leading children to believe the equals sign has an asymmetrical meaning such as “makes” or “gives the answer”. Interventions that introduce reversible and more symmetrical reasoning about equivalence in problem solving partially alleviates these difficulties (Johannes & Davenport, 2017; Johannes et al., 2021). Our findings suggest that instructors’ explicit use of asymmetrical linguistic terms like “make” rather than “is the same as” might be an additional cause for children not obtaining a symmetrical meaning of the symbol “ $=$ ”. If at school entry children already know which common English predicates are symmetrical and asymmetrical, pedagogically leveraging this understanding when arithmetic symbols are first introduced may avoid the problem entirely.

Concluding Remarks

The property of symmetry goes far beyond sensory experience, to social situations (e.g., *marry*, *meet*) and even to the abstractions pervasive in mathematical and scientific reasoning (e.g., *equal*, *similar*). Our findings support the existence of an abstract property of symmetry which humans access via both perceptual means (from certain observable events or situations) and linguistic means (from terms that token concepts with this property). More broadly, this work sheds light on the rich, structured nature of the language-cognition interface, and points towards a possible avenue for acquisition of word-to-world mappings for the seemingly inaccessible logical symmetry of linguistic terms.

Endnote

With great sadness, we wish to let readers know that our friend, collaborator, colleague, and mentor, Lila R. Gleitman, passed away while this project was nearing completion. It is not an overstatement to say that most of the ideas and motivation for this work came from Lila’s lifelong intellectual passion for the topic of symmetry, and especially, its manifestation in language. She was fully involved in earlier drafts of this paper and in presentations of this work at scientific conferences before her death. However, she was not able to participate in the final versions of the paper, which likely accounts for the absence of the characteristic flair and eloquence for which she was widely known. We, the remaining authors, have completed the paper always thinking of Lila, taking inspiration from our fond remembrances of her, and especially, her brilliant analyses of language and its relationship to cognition. We can only hope that she would have been pleased with the outcome.

Acknowledgments

We thank Jami and Dan Fisher (both of whom are bilingual ASL signers and English speakers) for their help in preparing the stimuli for Experiment 3. We thank Jenna Disefano and Abimael Hernandez Jimenez for assistance with stimulus preparation and data collection. We also thank Abby Clements, Chaz Firestone, Victor Gomes, Tyler Knowlton, and Barbara Partee for helpful discussion and/or comments. This work was supported by an NSF Collaborative Research Award (BCS-1941006 awarded to L.R.G. and J.C.T. at U. Penn and BCS-1941014 to B.L. at JHU), and NSF SBE Postdoctoral Research Fellowship SMA-2105228 awarded to A.H.

References

- Baylis, G. C., & Driver, J. (2001). Perception of symmetry and repetition within and across visual shapes: Part-descriptions and object-based attention. *Visual Cognition*, 8(2), 163–196.
- Bornstein, M. H., & Krinsky, S. J. (1985). Perception of symmetry in infancy: The salience of vertical symmetry and the perception of pattern wholes. *Journal of Experimental Child Psychology*, 39(1), 1–19.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3), 904–911.
- Cattaneo, Z., Fantino, M., Silvanto, J., Tinti, C., Pascual-Leone, A., & Vecchi, T. (2010). Symmetry perception in the blind. *Acta psychologica*, 134(3), 398–402.
- Cavanagh, P. (2021). The language of vision. *Perception*, 50(3), 195–215.
- Chestnut, E. K., & Markman, E. M. (2016). Are Horses Like Zebras, or Vice Versa? Children’s Sensitivity to the Asymmetries of Directional Comparisons. *Child Development*, 87(2), 568–582.
- Chomsky, N. (1981). *Lectures on Government and Binding: The Pisa Lectures*. Dordrecht: Foris.
- Dimitriadis, A. (2008). The event structure of irreducibly symmetric reciprocals. In J. Dölling, T. Heyde-Zybatow, & M. Schäfer (Eds.), *Event Structures in Linguistic Form and Interpretation*. De Gruyter.
- Fisher, C., Jin, K., & Scott, R. M. (2019). The developmental origins of Syntactic Bootstrapping. *Topics in Cognitive Science*, 12(1), 48–77. <https://doi.org/10.1111/tops.12447>
- Fodor, J. A. (1983). *The modularity of mind*. MIT press.
- Gattis, M. (2004). Mapping relational structure in spatial reasoning. *Cognitive Science*, 28(4), 589–610. <https://doi.org/10.1016/j.cogsci.2004.02.001>
- Gleitman, L. R. (1990). The Structural Sources of Verb Meanings. *Language Acquisition*, 1(1), 3–55.
- Gleitman, L. R., Gleitman, H., Miller, C., & Ostrin, R. (1996). Similar, and similar concepts. *Cognition*, 58(3), 321–376.
- Gleitman, L. R., & Partee, B. H. (in review). Understanding symmetrical predicates.
- Gleitman, L. R., Senghas, A., Flaherty, M., Coppola, M., & Goldin-Meadow, S. (2019). The emergence of the formal category “symmetry” in a new sign language. *Proceedings of the National Academy of Sciences*, 116(4), 11705–11711.
- Glenberg, A. M. (2015). Few believe the world is flat: How embodiment is changing the scientific understanding of cognition. *Canadian Journal of Experimental Psychology*, 69, 165–171. <http://dx.doi.org/10.1037/cep0000056>
- Goldwater, M. B., & Gentner, D. (2015). On the acquisition of abstract knowledge: Structural alignment and explication in learning causal system categories. *Cognition*, 137, 137–153. <https://doi.org/10.1016/j.cognition.2014.12.001>
- Goldwater, M. B., Markman, A. B., & Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition*, 118(3), 359–376. <https://doi.org/10.1016/j.cognition.2010.10.009>
- Goodman, N. (1955). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Hafri, A. (2022, June 28). Where word and world meet: Language and vision share an abstract representation of symmetry. Retrieved from <https://osf.io/64dnq>
- Hafri, A., Boger, T., & Firestone, C. (2022). Melting ice with your mind: Representational momentum for physical states. *Psychological Science*. <https://doi.org/10.1177/09567976211051744>
- Hafri, A., & Firestone, C. (2021). The perception of relations. *Trends in Cognitive Sciences*, 25, 475–492.
- Hafri, A., Gleitman, L.R., Landau, B., & Trueswell, J.C. (2021). When word and world meet: Intuitive correspondence between visual and linguistic symmetry. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*. Vienna, Austria.
- Hafri, A., Landau, B., Bonner, M.F., & Firestone, C. (2020). A phone in a basket looks like a knife in a cup:

- The perception of abstract relations. *PsyArXiv*. <https://doi.org/10.31234/osf.io/jx4yg>
- Hafri, A., Trueswell, J. C., & Strickland, B. (2018). Encoding of event roles from visual scenes is rapid, spontaneous, and interacts with higher-level visual processing. *Cognition*, 175, 36–52. <https://doi.org/10.1016/j.cognition.2018.02.011>
- Hochmann, J.-R., Tuerk, A. S., Sanborn, S., Zhu, R., Long, R., Dempster, M., & Carey, S. (2017). Children's representation of abstract relations in relational/array match-to-sample tasks. *Cognitive Psychology*, 99, 17–43.
- Jackendoff, R. (1987). On Beyond Zebra: The relation of linguistic and visual information. *Cognition*, 26(2), 89–114. [https://doi.org/10.1016/0010-0277\(87\)90026-6](https://doi.org/10.1016/0010-0277(87)90026-6)
- Johannes, K., & Davenport, J. (2017). Promoting Children's Relational Understanding of Equivalence. *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*.
- Johannes, K., McNeil, N., Kao, Y., & Davenport, J. (2021). Predicting Learning and Knowledge Transfer in Two Early Mathematical Equivalence Interventions. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society*.
- Kipper, K., Korhonen, A., Ryant, N., & Palmer, M. (2008). A large-scale classification of English verbs. *Language Resources and Evaluation*, 42(1), 21–40. <https://doi.org/10.1007/s10579-007-9048-2>
- Landau, B., & Gleitman, L. R. (1985). *Language and experience: Evidence from the blind child*. Harvard University Press.
- Landau, B., & Gleitman, L. R. (2015). Height Matters. In I. Toivonen, P. Csur, & E. van der Zee (Eds.), *Structures in the mind: Essays on language, music, and cognition in honor of Ray Jackendoff* (pp. 187–210). The MIT Press.
- Levin, B. (1993). *English Verb Classes and Alternations: A Preliminary Investigation*. Chicago, IL: University of Chicago Press.
- Mahon, B. Z. (2015a). The burden of embodied cognition. *Canadian Journal of Experimental Psychology*, 69, 172–178. <http://dx.doi.org/10.1037/cep0000060>
- Mahon, B. Z. (2015b). Response to Glenberg: Conceptual content does not constrain the representational format of concepts. *Canadian Journal of Experimental Psychology*, 69, 179–180. <http://dx.doi.org/10.1037/cep0000059>
- Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, 25(4), 431–467. <https://doi.org/10.1006/cogp.1993.1011>
- Markman, A. B., & Gentner, D. (2000). Structure mapping in the comparison process. *American Journal of Psychology*, 113(4), 501–538.
- McNeil, N. M. (2014). A “change-resistance” account of children's difficulties understanding mathematical equivalence. *Child Development Perspectives*, 8, 42–47.
- Meyerhoff, H. S., & Suzuki, S. (2018). Beep, be-, or -ep: The impact of auditory transients on perceived bouncing/streaming. *Journal of Experimental Psychology: Human Perception and Performance*, 44(12), 1995–2004. <https://doi.org/10.1037/xhp0000585>
- Michotte, A. (1946/1963). *The perception of causality*. New York, NY: Basic Books (Original work published 1946).
- Miller, C. A. (1998). It Takes Two To Tango: Understanding and Acquiring Symmetrical Verbs. *Journal of Psycholinguistic Research*, 27(3).
- National Council of Teachers of Mathematics. (2000). *Principles and standards for school mathematics*. Reston, VA.
- Palmer, S.E. and Hemenway, K. (1978). Orientation and symmetry: effects of multiple, rotational, and near symmetries. *Exp. Psychol Hum. Percept. Perform.* 4, 691–702
- Partee, B. H. (2008). Symmetry and symmetrical predicates. In A. E. Kibrik, V. I. Belikov and B. V. Dobrov & B. V. Dobrovolsky (eds.), *Computational linguistics and intellectual technologies: Papers from the international conference “Dialogue”*, 606–611. Moscow: Institut Problem Informatiki. Extended version appears as unpublished

- notes at <http://people.umass.edu/partee/Research.htm>, accessed: 10 May 2021.
- Partee, B. H. & Gleitman, L.R. (under review). Symmetry and Symmetrical Predicates.
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153–163. <https://doi.org/10.1016/j.jesp.2017.01.006>
- Quilty-Dunn, J. (2020) Concepts and predication from perception to cognition. *Philosophical Issues*, 30, 273–292.
- Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology*, 21(2), 139–155.
- Rolfs, M., Dambacher, M., & Cavanagh, P. (2013). Visual adaptation of the perception of causality. *Current Biology*, 23(3), 250–254. <https://doi.org/10.1016/j.cub.2012.12.017>
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80(1–2), 1–46.
- Sekuler, R., Sekuler, A. B., & Lau, R. (1997). Sound alters visual motion perception. *Nature*, 385(6614), 308–308. <https://doi.org/10.1038/385308a0>
- Shao, R., & Gentner, D. (2019). Symmetry: Low-level visual feature or abstract relation? *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*.
- Strickland, B. (2017). Language reflects “core” cognition: A new theory about the origin of cross-linguistic regularities. *Cognitive Science*, 41(1), 70–101. <https://doi.org/10.1111/cogs.12332>
- Strickland, B., Geraci, C., Chemla, E., Schlenker, P., Kelepir, M., & Pfau, R. (2015). Event representations constrain the structure of language: Sign language as a window into universally accessible linguistic biases. *Proceedings of the National Academy of Sciences*, 112(19), 5968–5973.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327–352.
- van Buren, B., Gao, T., & Scholl, B. J. (2017). What are the underlying units of perceived animacy? Chasing detection is intrinsically object-based. *Psychonomic Bulletin & Review*, 24(5), 1604–1610.
- Wagemans, J. (1997). Characteristics and models of human symmetry detection. *Trends in Cognitive Sciences*, 1(9), 346–352.
- Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological Bulletin*, 138(6), 1172–1217.
- Warriner, A.B., Kuperman, V. & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavioral Research Methods*, 45, 1191–1207 (2013). <https://doi.org/10.3758/s13428-012-0314-x>
- Wege, T. E., Batchelor, S., Inglis, M., Mistry, H., & Schlimm, D. (2020). Iconicity in mathematical notation: Commutativity and symmetry. *Journal of Numerical Cognition*, 6(3), 378–392.
- Weyl, H. (1952). *Symmetry*. Princeton University Press. Princeton, New Jersey.
- Williams, A. (2021). Events in Semantics. In P. Stalmaszczyk (Ed.), *The Cambridge Handbook of the Philosophy of Language* (Cambridge Handbooks in Language and Linguistics, pp. 366–386). Cambridge: Cambridge University Press. doi:10.1017/9781108698283.021
- Winter, Y. (2018). Symmetric predicates and the semantics of reciprocal alternations. *Semantics and Pragmatics*, 11, Article 1.
- Yuan, S., & Fisher, C. (2009). “Really? She blinked the baby?”: two-year-olds learn combinatorial facts about verbs by listening. *Psychological Science*, 20(5), 619–626. <https://doi.org/10.1111/j.1467-9280.2009.02341.x>

Appendix A

The table below shows the frequency (log base 10), concreteness, and valence (“pleasantness”) of each English predicate used in our studies. Concreteness and frequency (from SUBTLEX) were from the norms in Brysbaert et al. (2014); valence was from the norms in Warriner et al. (2013). The concreteness scale ranged from 1 to 5, and valence from 1 to 9. For most predicates, we used the stem form of each word (e.g., “marry”, “collide”) to extract these norms. For “box”, “match”, and “lead”, we used the progressive form of the predicates, as the noun homographs of these predicates are quite common.

Note that symmetrical predicates and their non-symmetrical foils were closely matched on most of these properties, as well as on word length, such that only valence showed a marginally significant difference between the two paired groups (word length: $t(23) = 0.21$, $p = .833$; log frequency: $t(23) = 0.07$, $p = .946$; concreteness: $t(23) = 1.22$, $p = .234$; valence: $t(23) = 1.95$, $p = .063$).

Symmetrical predicates are ordered in this table according to their mean symmetry rating. Non-symmetrical predicates are sorted according to the order in which their yoked symmetrical counterparts appear (i.e., by the “pair number”).

Pair Number	Type	Predicate	Frequency (Log 10)	Concreteness	Valence	Mean Symmetry Rating
1	Sym	be identical	2.452	3.5	5.3	5.92
2	Sym	be similar	2.843	2.46	5.86	5.88
3	Sym	marry	3.726	3.03	7.09	5.72
4	Sym	date	3.858	3.9	7.18	5.6
5	Sym	match	3.402	2.82	6.3	5.6
6	Sym	collide	1.681	4	3.1	5.56
7	Sym	collaborate	1.505	1.81	6.15	5.56
8	Sym	tango	2.439	4.1	5.47	5.4
9	Sym	chat	2.92	3.66	5.75	5.36
10	Sym	meet	4.255	3	6.09	5.36
11	Sym	equal	2.834	2.56	6.47	5.32
12	Sym	interact	1.869	2.88	6	5.24
13	Sym	unite	2.19	2.58	6.4	5.24
14	Sym	agree	3.518	2.31	7.17	5.12
15	Sym	combine	2.057	3.41	4.84	5
16	Sym	debate	2.677	3.18	4.95	5

Pair Number	Type	Predicate	Frequency (Log 10)	Concreteness	Valence	Mean Symmetry Rating
17	Sym	intersect	1.23	3.62	3	4.84
18	Sym	negotiate	2.535	2	5.57	4.56
19	Sym	clash	1.839	2.63	3.96	4.52
20	Sym	separate	3.045	3.25	3.95	4.4
21	Sym	correspond	1.69	2.27	5.75	4.24
22	Sym	disagree	2.53	2.77	2.84	4.2
23	Sym	box	3.661	4.55	4.47	3.72
24	Sym	differ	2.097	1.59	5.1	3.72
1	Non-Sym	be inferior	2.196	1.7	3.43	1.28
2	Non-Sym	be typical	2.808	1.52	4.58	3.6
3	Non-Sym	adopt	2.378	2.36	6.9	2.28
4	Non-Sym	befriend	1.362	2.36	7.14	4.68
5	Non-Sym	gauge	2.053	4	5.17	2.68
6	Non-Sym	hit	4.147	4.11	3.95	2.52
7	Non-Sym	contribute	2.29	2.13	6.15	3.12
8	Non-Sym	lead	3.628	3.14	5.56	1.68
9	Non-Sym	tell	4.944	2.9	5.27	2.12
10	Non-Sym	greet	2.43	2.96	6.25	4.08
11	Non-Sym	exceed	1.813	1.76	4.79	1.6
12	Non-Sym	intervene	1.863	2.57	5.21	2.64
13	Non-Sym	dominate	1.982	2.12	4.4	1.28
14	Non-Sym	consent	2.58	2.08	6.42	3.88
15	Non-Sym	expand	2.253	3.36	5.35	2.72
16	Non-Sym	lecture	2.728	3.79	4	1.8
17	Non-Sym	interfere	2.688	2.59	3.94	2.48

Pair Number	Type	Predicate	Frequency (Log 10)	Concreteness	Valence	Mean Symmetry Rating
18	Non-Sym	propose	2.822	2.19	6.26	1.92
19	Non-Sym	confront	2.34	2.43	4.59	2.68
20	Non-Sym	withdraw	2.529	3.04	4.11	2.44
21	Non-Sym	contact	3.519	3.86	6.17	3.88
22	Non-Sym	reject	2.337	2.4	2.95	2.16
23	Non-Sym	punch	3.18	4.39	3.27	2.56
24	Non-Sym	alter	2.389	3.07	4.57	2.64