Real-world size is automatically encoded in preschoolers' object representations

Bria Long^{1,2}, Mariko Moher³, Susan Carey¹ & Talia Konkle¹, ¹Department of Psychology, Harvard University ²Department of Psychology, Stanford University ³Department of Psychology, Williams College

Corresponding author: Bria Long, <u>bria@stanford.edu</u> Jordan Hall, 450 Serra Mall, Stanford, CA 94305

Abstract

When adults see a picture of an object, they automatically process how big the object typically is in the real world (Konkle & Oliva, 2012a). How much life experience is needed for this automatic size processing to emerge? Here, we ask whether preschoolers show this same signature of automatic size processing. We showed 3- and 4-year-olds displays with two pictures of objects and asked them to touch the picture that was smaller on the screen. Critically, the relative visual sizes of the objects could either be congruent with their relative real-world sizes (e.g., a small picture of a shoe next to a big picture of a car) or incongruent with their relative real-world sizes (e.g., a big picture of a shoe next to a small picture of a car). Across two experiments, we found that preschoolers were worse at making visual size judgments on incongruent trials, suggesting that real-world size was automatically activated and interfered with their performance. In a third experiment, we found that both 4-year-olds and adults showed similar item-pair effects (i.e., showed larger Size-Stroop effects for the pairs of items, relative to other pairs). Furthermore, the magnitude of the item-pair Stroop effects in 4-year-olds did not depend on whether they could recognize the pictured objects, suggesting that the perceptual features of these objects were sufficient to trigger the processing of real-world size information. These results indicate that, by 3-4 years of age, children automatically extract real-world size information from depicted objects.

Research Highlights

- Real-world size interfered with 4-year-olds' ability to make visual size judgments about pictured objects in the Size-Stroop task
- The same pairs of objects generated robust Size-Stroop effects in both adults and 4year-olds
- This was true even when 4-year-olds could not name the pictured objects

Keywords:

Object representation; real-world size; Stroop effects; Visual development

When we look at the world, we easily recognize objects and perceive their physical size, from small objects like cups and paperclips, to bigger objects like cars and pianos. Indeed, our representations of the typical sizes of objects enter into human mental life in many ways—for example, providing the standards for the application of words like "big" and "small" (a small car is smaller than average for cars, but nonetheless much larger than a large cup), participating in computations of spatial layout (in their role in specifying objects' distance from us), and constraining motor interactions (e.g. we tend to pick up small objects with our hands, and we need to navigate around big objects). Thus, the real-world size of objects is likely a primary factor that structures our visual experience with objects as we learn about the world.

And indeed, there is evidence that real-world size has an organizing role in both perceptual and neural object representation by adulthood (Henik et al., 2017; Julian, Ryan, & Epstein, 2016; Konkle & Oliva, 2012a; Konkle & Oliva 2012b). For example, even though there are many kinds of big and small objects, visual search behavior shows that big objects, as a class, actually *look* different than small objects, as a class (Long et al., 2016). Furthermore, at a neural level, the distinction between small vs. big objects also organizes responses in occipitotemporal cortex (Cate et al., 2011; Konkle & Oliva, 2012b; Khaligh-Razavi et al., 2018; Julian, Ryan, & Epstein, 2016).

In fact, real-world size information appears to be so ingrained in adult object representations that when we see an object we not only automatically recognize what it is, but automatically activate knowledge about its typical size in the real world (e.g., Chiou & Ralph, 2016; Gliksman et al., 2016; Konkle & Oliva, 2012b; Sellaro et al., 2009; see also Paivio, 1975). As evidence of this, in a Size-Stroop paradigm, adults were asked to make a *visual* size judgment about which of two images is bigger (or smaller) on the screen, while

ignoring the objects' sizes in the real world. Critically, adults were slower and less accurate at making visual size judgments when the images' relative visual sizes were incongruent with the relative real-world sizes of the depicted objects (i.e., a big picture of a teapot and a small picture of a gazebo) versus when they were congruent with their real-world sizes (i.e., a big picture of a gazebo and a small picture of a teapot; Konkle & Oliva, 2012a; see Appendix Figure S1). Thus, in this task, even though real-world size information was taskirrelevant, it was automatically activated and interfered with adults' ability to make visual size judgments.

It is intuitive to assume that one must recognize an object before any real-world size related processing can proceed, as in classic accounts of semantic processing (e.g., Collins & Quillian, 1969; Jolicoeur, Gluck, & Kosslyn, 1984). However, we have shown that realworld size information can be activated without first recognizing the object, using "texform" stimuli. Texform stimuli preserve mid-level features like curvature and texture but lack the high-level features that enable basic-level recognition (see Appendix, Supplementary Figure 1). In a line of behavioral and neural studies with adults, we have found that when images are transformed into texforms, (1) visual search effects between big and small objects persist (Long et al., 2016), as do (2) neural differences between big and small object images (Long, Yu, & Konkle, in press), as well as (3) the Size-Stroop effect (Long & Konkle, 2017). These studies demonstrate that there are *mid-level perceptual* features that systematically distinguish small objects vs. big objects as classes, that underlie object representations along the ventral stream, and that can lead to the automatic computation of real-world size.

The fact that real-world size is such an ingrained and organizing property of our object representations raises several critical developmental questions: how does this real-

world size sensitivity emerge over development, i.e., what, if any, innate support exists for it, and what learning mechanisms are involved? Answering these questions is not only an important project within developmental cognitive neuroscience, but also would shed light on *how* the distinction between small vs. large objects as classes becomes an organizing property of adults' object representations. Here we take a first step in this developmental project, starting with preschoolers, as they are the youngest age group likely capable of performing the same exact tasks used to study object size processing in adults.

Preschoolers are at an interesting point in the development of their object representations. Early in the preschool years, by age 2, children can say when an object is "big" or "little" with respect to other objects of the same kind (e.g., mittens), indicating that they explicitly represent the average sizes of some categories (Ebeling & Gelman, 1988; Gelman & Ebeling, 1989). Furthermore, evidence from visual search suggests that big objects as a class "look" different than small objects as a class to children as young as age 3 (Long, Moher, Carey & Konkle, under review), indicating that real-world object size influences preschoolers' perceptual similarity computations. At the same time, however, considerable evidence suggests that preschoolers' object representations may be far from mature (for reviews, see Nishimura et al., 2009; Jüttner et al., 2016). For example, children continue to exhibit deficits recognizing objects across wide variations in lighting and pose throughout middle childhood (e.g., Bova et al., 2007), only integrate haptic and visual information in size discrimination tasks around 8 years of age (Gori et al., 2008), and are less deceived by the Ebbinghaus illusion than adults until 7 years of age (Doherty et al., 2010). Further, as children's *own* physical size relative to objects changes dramatically over the first few years of their life, one might expect their object size representations to mature gradually throughout childhood.

Thus, the present experiments seek to establish whether preschoolers, like adults, automatically activate real-world size information when they see pictured objects. To do so, in Experiments 1 & 2, we use the Size-Stroop task to investigate whether 3- and 4-year-old children, like adults, automatically activate the real-world size of pictured objects, even when this information interferes with the task at hand. Further, and more speculatively, we aim to assess whether the data support mid-level perceptual processing vs. basic-level recognition as a locus of the observed effects. In Experiment 3, we explore which item pairs generate the greatest Size-Stroop effects as a first step towards understanding if similar perceptual mechanisms underlie real-world size representations in preschoolers and adults.

Experiment 1: Do preschoolers show the Size-Stroop effect?

We adapted the Size-Stroop task (Konkle & Oliva, 2012a) for children by converting it to an iPad game. Children were asked to "touch the picture that is smaller on the screen." Critically, the identity of the objects and their real-world sizes are completely irrelevant to the task. However, if preschoolers automatically activate information about objects' typical sizes in the real world during this task, then they should be slower and less accurate on incongruent displays, in which the object that is bigger in the real world is smaller on the screen, than on congruent displays, in which the object that is bigger in the real world is bigger on the screen.

Methods

Participants. Eighty 3- and 4-year-old children participated, either at the Boston Children's Museum, the Harvard Lab for Developmental Studies, or the Williams College Children's Center. A parent gave consent prior to participation, and the Institutional Review Board at Harvard University approved the study. We aimed to recruit enough

participants to include approximately double the number of subjects needed to observe the effect in adults in each age group (N = 16 in adults; see Konkle & Oliva, 2012a). One child began the task but did not complete more than two trials. This left us with 79 children in the final sample, with 48 three-year-olds (M = 41.8 months, SD = 3.0 months) and 31 four-year-olds (M = 53.7 months, SD = 3.4 months).



Figure 1. Example stimuli for Experiments 1 and 2. In congruent displays, the relative size of the objects were consistent with their sizes in the real world, and in incongruent displays, the relative size of the two objects were inconsistent with their size in the real world.

Experimental Set-Up. Children sat at a table across from an experimenter who held an iPad for them (see Figure 2). The experimenter could not see the images on the screen and was thus blind to condition. Experiments were run on an iPad in a web-browser (Safari) and custom code was written in Javascript using the JQuery toolbox. Reaction time, touch position, accuracy, and experimental details were recorded and saved after each trial to an online database.

Stimuli. Images were identical to those used in Experiment 1B of Konkle & Oliva (2012a); these images of 20 big objects and 20 small objects were matched in terms of their overall area and paired by their vertical height. The same pairs of big and small objects were always presented together on both congruent and incongruent trials (see Figure 1).

Procedure. There were two phases to the experiment. First, practice trials verified that the child could make visual size judgments about geometric shapes. Next, there was a test phase where children made visual size judgments about two pictured objects.

The first 35 out of 80 children received a paper version of the practice phase. These children were presented with two different colored shapes, one of which was bigger than the other, and were asked to "Touch the shape that is smaller *on this paper*." All 35 children completed five correct practice trials. However, as several children were distracted by the appearance of the iPad for the test phase, the remaining children completed the practice phase on the iPad. These 45 children completed nine correct practice trials before the test phase.¹ Here, children touched a blue dot to begin each trial, after which they were presented with two different colored shapes, one of which was bigger than the other. Children were asked to "Touch the shape that is smaller *on the screen*." These last three words ("on the screen") were emphasized to clarify any ambiguity in these instructions. When children answered correctly, the iPad played a pleasant sound and advanced to the next practice trial. The experimenter also reinforced on-task performance by saying "good job!" when children selected the correct target.

¹ Children in these two familiarization versions did not perform more or less accurately on test trials (no main effect of familiarization version on error rates; F(1,74) = 3.35, p = 0.07) and or on congruent versus incongruent displays (no interaction of familiarization version with trial type on error rates; F(1,74) = .62, p = 0.43).



Figure 2. Schematic of practice and test trials used in Experiments 1 and 2. A child performing the task is shown on the right.

In the test phase, at the beginning of each trial, all children were asked to "touch the blue dot to begin." After children touched the blue dot, there was a brief delay of 500ms after which two images appeared on either side of the screen. Children were asked to "Touch the picture that is smaller *on the screen*."² Critically, there were two different kinds of trials: congruent trials, when the relative real-world sizes of the picture objects were congruent with their relative visual sizes on the screen (i.e., a big picture of a car and a small picture of a cup) and incongruent trials, when the relative real-world sizes on the screen (i.e., a small picture of a car and a big picture of a cup). If the child selected the correct image, a pleasant sound was played; if the child selected the incorrect image, no sound was played. In either case, the blue dot then reappeared to signal the start of the next trial. To encourage accuracy, a picture of Mickey Mouse also appeared after every 3 correct trials,

² In the adult study on which this study is based (Konkle & Oliva, 2012a), half of the participants were asked to indicate which object is larger on the screen, and half of the participants were asked to indicate which object is smaller on the screen. Size-Stroop effects were observed in reaction times and error rates for both tasks. However, the *indicate-smaller* task produced a slightly bigger effect size, and thus, to maximize power, children were only asked the latter question.

and children's progress was marked with a stamp by the experimenter.³ The experimenter also periodically gave positive feedback, saying "good job!," noting how many stamps the child had acquired, and encouraging children to keep playing the game. Children continued until they completed a maximum of 80 trials or wanted to stop the experiment.

Counterbalancing. Each pair of big and small objects appeared in both incongruent and congruent configurations. In addition, the visually bigger object appeared on both sides of the screen, creating 4 displays per pair of objects and 80 total possible displays. Every combination of target side (right, left) and trial type (congruent, incongruent) appeared every four trials during the experiment. The image pair that occurred on a given trial was randomized throughout each session for each child.

Data Inclusion. We assessed both accuracy and reaction time as dependent measures, both across all participants and separately for 3- and 4-year-olds. To do so, we adopted the following exclusion criteria and data-trimming methods. First, we excluded all geometric shape practice trials and, a priori, the first 10 trials from the test phase.⁴ This yielded an average of 51.7 trials out of a possible 70 (range 4 to 70), in 79 children, which were used in the accuracy analysis. Three children (all 3-year-olds) did not complete more than 5 trials in each condition after these first 10 trials and were excluded from all subsequent analyses. Error analyses were thus conducted on the remaining 76 children.

For reaction time analyses, we additionally excluded incorrect trials and trials with RTs slower than 4 seconds, a preset criterion (6.6% of correct trials). This RT cutoff has previously been used as a cutoff when analyzing preschooler's reaction times in a

³ In a pilot study, we found that marking children's progress on the stamp sheet dramatically increased the number of trials children were willing to complete, suggesting that children were very sensitive to this feedback.

⁴ This cutoff of 10 trials was chosen after piloting the task in a separate group of children and noticing that some children were still responding very slowly during the first few trials.

touchscreen-based task (Frank et al., 2016), in order to eliminate extra-long trials where children are likely off-task. Children were included if, after this RT trimming procedure, they had at least 5 correct trials per condition (congruent, incongruent) excluding practice trials. Four additional children were excluded for not meeting these criteria, who were all 3-year-olds. This left us with 72 children for RT analyses: 41 three-year-olds (M = 42.2 months, SD = 2.9 months) and all of the 31 four-year-olds. On average, 3-year-olds contributed 47.4 trials to RT analyses, and 4-year-olds contributed 47.8 trials to RT analyses.

Data Analysis. We analyzed error patterns and reaction times in two ways. First, we performed an ANOVA to assess the effects of the within-subjects variable of trial type (congruent versus incongruent) and the between-subject variable of age group (3 versus 4 vears) on error rates (percentage of completed trials that were errors) and reaction times. Post-hoc tests are reported using one-tailed t-tests, as the results are only interpretable if children performed worse on incongruent relative to congruent displays. Second, to ensure the robustness of our results, we modeled log-transformed RT data in a linear mixed effect model (as RT data is non-normally distributed, Whelan, 2008) and modeled error patterns using a generalized linear mixed effect model. Models were fit with a linear mixed-effect model code package implemented in R (lme4), where age group, congruency, and their interaction were specified as fixed effects. Random intercepts for subjects and for individual displays were always included, and random slopes were included if the model was able to converge with this more maximal design (Barr et al., 2013). With these random effects terms, the statistical models are better able to (i) account for the different numbers of trials completed by individual children, and also (ii) to ensure that the results are not

strongly driven by particular displays. All data and analysis code are available at the public repository for this manuscript (<u>www.github.com/brialorelle/kidstroop</u>).

Results

Error Results. Children made relatively few errors (M = 10.6%) suggesting they understood the task instructions, though 3-year-olds made more errors than 4-year-olds (main effect of age: 3-year-olds: M = 14.1%, 4-year-olds: M = 5.6%, F(1,74) = 7.08, p = 0.01, η G2 = 0.07).



Figure 3: Average error rates (upper panel) and reaction times (lower panel) are shown for congruent (blue/light) and incongruent (red/dark) displays for 3-year-olds and 4-year-olds in Experiment 1 (A) and from 4-year-olds in Experiment 2 (B, replication experiment). Error bars represent bootstrapped 95% confidence intervals.

Critically, we found that children showed evidence for the Size-Stroop effect in their errors; they made more errors on incongruent than congruent displays (main effect of trial type: congruent M = 8.0% (SD = 12.1%), incongruent M = 13.2% (SD = 18.3%), F(1,74) = 11.87, p < .001, η G2 = 0.03). The Size-Stroop effect was apparent throughout this age range;

there was no interaction between age group and trial type (F(1,74) = 0.31, p = 0.58, $\eta G2 < 0$). Further, planned ad-hoc comparisons confirmed that the Size-Stroop effect was observed at *each* age: 3-year-olds: congruent M = 11.2%, incongruent M = 17.0%, t(44) = 2.88, p = 0.01; 4-year-olds: congruent M = 3.5%, incongruent M = 7.7%, t(30) = 2.2, p = 0.02; see Figure 3A. The generalized mixed effect model confirmed these analyses while accounting for variance across displays and subjects (main effect of congruency; B = 0.56, SE = 0.12, Z = 4.89, p < .001).

Reaction Time Results. Considering both 3- and 4-year-olds together, we found that children did not take longer to make visual size judgments on incongruent versus congruent displays (no main effect of trial type: congruent M= 1758ms, incongruent M = 1778ms, F(1, 70) = 0.9, p = 0.35, η G2 < 0.001). However, considering 3- and 4-year-olds separately, we found that 4-year-olds showed the Size-Stroop effect in their RTs (congruent M = 1555 ms, SD = 359 ms, incongruent M = 1622 ms, SD = 319 ms, t(30) = 2.37, p = 0.01, Cohen's d = 0.43), while the 3-year-olds did not (congruent M = 192 ms, SD = 475 ms, incongruent M = 1901 ms, SD = 446 ms, t(40) = -0.54, p = 0.7, Cohen's d = -0.08, Figure 3A). This same pattern of results was evident in the linear mixed effect model: that is, when combining across all children there was no congruency effect in reaction time (B = 0.048, SE = 0.025, t = 1.9, p = 0.062), but congruency was significant when 4-year-olds were considered separately (congruency, B = 0.047, SE = 0.021, t = 2.21, p = .04).

As is evident in Figure 3A, the 3-year-olds generally took longer to respond on the iPad than 4-year-olds. Thus, we conducted an exploratory analysis to examine whether age or overall slowness was more likely to account for the 3-year-olds' lack of the Size-Stroop effect on RTs. First, we analyzed whether children's age (in months) was correlated with the degree to which children made more errors or had slower RTs on the incongruent than

the congruent trials. Age was only weakly correlated with the size of the Size-Stroop effect for RTs (RTs: r = .20, p = .09) and was not positively correlated for the size of the Size-Stroop effect on errors (Error rates: r = .10, p = .22). Next, we analyzed how overall reaction time was related to Stroop RT and Stroop errors. We found that children who performed the task more slowly also tended to show either a very positive or a very negative Size-Stroop RT effect (average RT correlated with positive Stroop effects, r = .23, p = .06; average RT correlated with absolute valued Stroop effects, r = 0.33, p < .01; see Supplementary Figure 2). In other words, children who had slower reaction times also tended to have more variance in their RTs, leading to noisier estimates of the Size-Stroop effect.

Experiment 2: Replication

In Experiment 1, we found that both 3 and 4-year-olds showed a Size-Stroop effect in their error patterns, suggesting that preschoolers automatically activated information about the real-world sizes of the depicted objects. Further, we found that the Size-Stroop effect was also evident in 4-year-olds reaction times, as it in adults (Konkle & Oliva, 2012a). However, as 3-year-olds also tended to stay less on task and did not consistently make *speeded* visual size judgments, making harder to obtain accurate estimates of 3-year-olds' reaction times for congruent versus incongruent displays and thus to observe a Size-Stroop effect in their reaction times.

In Experiment 2 we sought to replicate the results of Experiment 1 in an independent group of 4-year-olds. Given that we initially hypothesized to see a reaction time effect in both age groups, and only observed it in the 4-year-olds, the aim of this experiment was to validate the robustness of this RT effect in 4-year-olds. Further, in Experiment 2, we a priori decided to add an additional exclusion criterion for overall slow

responders, as in Experiment 1 we found that children with very slow RTs tended to show highly variable Size-Stroop effects in their RTs.

Methods

Participants. Thirty-five 4-year-olds were recruited for Experiment 2 so that approximately the same number of 4-year-olds would contribute to RT analyses as in Experiment 1. Children were recruited and participated at the Boston Children's Museum or the Harvard Lab for Developmental Studies. One child began the task but did not complete more than two trials and was excluded from analysis. One other child participated but was excluded for parental interference, leaving us with 33 4-year-olds in our final sample (M = 53.3 months, SD = 3.2 months, 15 males).

Experimental Setup, Stimuli, & Counterbalancing. All aspects of Experiment 2 were identical to those of Experiment 1, except that all children did all practice trials on the iPad, and we encouraged children to obtain 20 stamps (rather than just as many stamps as possible), to help maximize the number of children who could be included in RT analyses.

Data inclusion. As in Experiment 1, the first 10 test trials were discarded for all participants. Children completed an average of 47.0 trials (range=28 to 56) out of a possible 70. All children were included in error analyses as they completed more than 5 trials in each condition after these first ten trials (range=14-29 trials per condition). For reaction time analyses, we applied the same exclusion criteria as in Experiment 1, excluded trials where children responded incorrectly or took longer than 4 seconds to respond (M = 1.2% of correct trials). No children were excluded on the basis of not having 5 or more test trials with correct responses made in less than 4 seconds. As planned, we then excluded children whose average RTs (across both conditions) were slower than 2 standard deviations from the average group RT (only 2 participants; mean RTs=2603 ms, 2432 ms,

z-scores=3.1, 2.5). After applying these inclusion criteria, we analyzed the RTs of 31 children, who completed an average of 45.26 trials after practice.

Analysis. The exact same analysis plan (ANOVAs, linear mixed effect modeling) from Experiment 1 was followed in Experiment 2.

Results & Discussion

Error Results. Overall, we replicated the main finding from Experiment 1, finding that 4-year-olds made more errors on incongruent displays than congruent displays (incongruent M = 3.7%, congruent M = 1.4%, t(32) = 2.55, p = .008), even though they made fewer errors overall when compared to 4-year-olds in Experiment 1 (see Figure 3B). This result was confirmed in our mixed effect model (main effect of congruency, B = 1.216, SE = 0.56, t = 2.19, p = 0.028)

Reaction Time Results. Further, 4-year-olds also exhibited a Size-Strop effect in their reaction times, taking longer to make visual size judgments on incongruent displays versus congruent displays (congruent M = 1438 ms, incongruent M = 1480 ms, t(30) = 2.30, p = .01, Cohen's d = .41, Figure 3B). Our linear mixed-effect model on logged RTs revealed the same pattern of results (B = .04, SE = .02, t = 2.28, p = .03).⁵ Thus, these data replicate the pattern of effects seen in Experiment 1; 4-year-olds exhibit a Size-Stroop effect in both their errors and reaction times.

Experiment 3: Size-Stroop Item-Pair Effects

Experiments 1 and 2 show that the Size-Stroop effect is observable in error rates by age 3 and in RTs by age 4, and thus establish that when preschool aged children see

⁵ As an exploratory analysis, we included the two children with slow overall RTs. We found that including these children did not change the pattern of effects in the linear mixed-effect model on log RTs (B = .03, SE = 0.1, t = 2.11, p = .04) but did change the pattern of effects in a traditional paired t-test (t(32) = 1.05, p = 0.15).

pictured objects, information about the real-world size of these objects is automatically activated. However, these results leave open the exact representations and computations underlying children's specification of real-word size of the pictured objects, and whether these are the same as those of adults. In the following post-hoc analyses, we begin to provide some insight into these questions by taking advantage of the fact that children and adults completed this task using the same exact images (Experiment 1 and 2 for children; Konkle & Oliva 2012a for adults). Further, for both adults and children there is variability in the magnitude of the Size-Stroop effect that a given item-pair generates; some item-pairs generate greater incongruency costs than do others. By comparing item-pair effects between children and adults, we can indirectly examine whether similar mechanisms underlie the Size-Stroop effect for preschoolers and for adults.

One possibility is that Size-Stroop item-pair effects are highly correlated between children and adults, which would provide indirect evidence that the mechanisms underlying these effects are similar. In adults, our prior research has highlighted that itempair effects elicited by texform stimuli are related to differences in mid-level visual features. For example, big objects tend to be boxy and small objects tend to be curvy, and these kinds of differences at this perceptual level are sufficient to trigger a Size-Stroop effect, even when the items on the display are not recognizable at the basic-level (Long & Konkle, 2017). If children are sensitive to the same mid-level visual features—that is, if children have abstracted the same mid-level features that characterize small objects and big objects as classes—then the same item-pairs should generate stronger Size-Stroopeffects in both children and adults. Our first analysis thus analyzes the degree to which children and adults' Stroop display effects are positively correlated.

However, it is also possible that children automatically compute the size of the depicted objects in this paradigm by first recognizing them as a kind (e.g. "cup"), and then retrieving information about the average size of that kind (e.g., "cups are small enough to be held with one hand"). To examine this possibility, we took advantage of the fact that the pictures of big and small objects used in this study were drawn from an adult study, and thus not all of the objects were necessarily recognizable by preschool aged children. Specifically, we analyzed whether in Experiments 1 and 2, the more recognizable objects (to preschoolers) give rise to a stronger Size-Stroop effect than do the less recognizable objects. For this analysis, we first asked an independent sample of preschool children to identify the pictured objects. Given the expected variation in recognizability, we next assessed whether the recognizability of the objects affected the magnitude of Stroop effects across item-pairs. If children's Size-Stroop effects rely on them first recognizing the objects, then we would expect to find stronger effects for pairs of items depicting more identifiable objects than for pairs of items depicting less identifiable objects. Our second analysis tested this hypothesis.

Analysis 1: Item-Pair Effects in Adults vs. Children

For our first analysis, we correlated the item-pair effects observed in the studies with preschool children and with adults. In both the present experiment and the original experiment with adults, stimuli were presented in consistent pairs; for example, a picture of a grill was always paired with a picture of a die on both incongruent and congruent trials. Thus, we could obtain measures of the Size-Stroop effect (RT to incongruent pair – RT to congruent pair) for each individual pair of objects for both children and adults. We used RT data from the 4-year-olds who contributed RT data in Experiment 1 and 2 to

calculate item-pair effects for 4-year-olds and the original data from Experiment 1B of Konkle & Oliva (2012a) to calculate item-pair effects for adults.

Methods. To calculate Stroop item-pair effects for children, we computed the average RTs for each congruent and incongruent display in each 4-year-old who contributed to RT analyses in either Experiments 1 or 2 (62 children). Next, we computed the average congruent and incongruent RTs for all 20 displays at the group level. Finally, Size-Stroop item-pair effects were calculated by subtracting the average congruent from incongruent RT at the group level, for each item-pair. The same procedure was used for the data from Experiment 1B of Konkle & Oliva (2012a). It is important to note that the procedure for this adult experiment is nearly exactly the same as that of the present Experiments 1 and 2. All of the changes that were made to this task (i.e., converting the task for use on a touchscreen interface, the appearance of Mickey Mouse every 3 correct trials) were for the purposes of adapting this task for use with children. Thus, the item-pair effects from this adult experiment are directly comparable with the item-pair effects obtained from Experiments 1 and 2.



Figure 4. Size-Stroop item-pair effects (Incongruent – Congruent RT) are plotted for each pair of objects for all 4year-olds in Experiment 1 and 2 as a function of adult's Stroop effects for the same pairs of displays. Shaded regions represent 95% confidence intervals around the mean. Pairs of displays are colored according to whether preschoolers were able to recognize both images in a pair were recognized greater than 75% of the time, based on the Object Identification experiment described in Analysis 2.

Results. Item-pair effects for preschoolers and adults were highly correlated (r = .64, p = .001; Figure 4); the same pairs of objects generated greater RT differences between congruent and incongruent item-pairs in both adults and children. Thus, these results provide indirect support that the same mechanisms underlie the Size-Stroop effect in adults and preschoolers and are consistent with the hypothesis that children may automatically compute real-world size from the same mid-level perceptual features adults rely upon to distinguish big from small objects as classes.

Analysis 2: Object Identification

We now turn to an analysis of the item-pair effects that explores whether preschoolers might also draw on knowledge of object size derived from basic level kind recognition. We first established the degree to which preschoolers could recognize the basic-level kind of each depicted object. Next, we assessed the degree to which variability in recognizability predicted the magnitude of the Size-Stroop effects.

Methods. 4-year-olds (*N* = 24) participated in the basic-level recognition task. Two additional children participated but were excluded because of (1) a speech articulation difficulty or (2) difficulty speaking English. Each child saw all 40 objects from Experiments 1 & 2 and was asked: "What does this look like?" If children did not provide a response, they were prompted with a broader question, "Does it remind you of anything?" and encouraged to guess. This second question was designed to elicit descriptions from children that could indicate whether they recognized the pictured object (even if they could

not name the object). Images were ordered such that no more than two items from the same size category appeared back to back.

Next, we coded all of these responses for *any* evidence of basic-level kind recognition. Some responses were straight forwardly correct (e.g. "apple", "desk"). However, we also counted descriptions of the object kind as correct (e.g. for a die, "you roll it and it gives you a number for a game"; see Supplemental Figure 3 for example responses from children and how they were coded). This more liberal method was followed because the aim here is separating objects that were identifiable to children from those objects that were not, regardless of whether the children knew the exact label. Overall, children identified the correct basic-level category of the objects 76.1% of the time, gave an incorrect answer 16.8% of the time, and did not give a response 7.1% of the time. Some items were always identified correctly (i.e., apple, 100% identification rate), while others were rarely identified correctly (i.e., perfume bottle, 33.3% identification rate).

Results. For the critical analysis, we separated displays into two groups: (1) displays where the basic-level identities of *both* the big and small depicted objects were well identified (greater than 75%, 8/20 pairs, M = 95.0% across all 16 items); (2) displays where one or more of the depicted objects was poorly identified (with recognition rates at 75% or less; 12/20 pairs, M = 63.54% across all 24 items). We found that pairs of objects that were well-identified at the basic level did not generate larger Size-Stroop effects in RTs (M = -5.4 ms) than pairs of objects that were not both well-identified (M = 85.2 ms;); if anything, we found a trend in the opposite direction (unpaired two-sample t-test, t(18) = -1.77, p = 0.09).⁶ For example, the Size-Stroop RT effect for the poorly recognized

⁶ As 75% was a relatively arbitrary cutoff, we examined whether a range of other cutoffs generated the same patterns of effects. Regardless of the cutoff we used, poorly- identified

barbecue/die pair was 66 ms, whereas the Size-Stroop RT effect for the well-recognized desk/apple pair was -120 ms. These results provide some evidence that children's Stroop effects are not driven by first recognizing the object, and then accessing the real-world size. However, the strength of this evidence is tempered by the fact that there are only a small number of item-pairs (*N*=20), and these were not specifically designed to probe for the role of explicit object recognition on the Size-Stroop effect.

In the next analysis, we considered the possibility that children might not recognize an object or describe its function but might systematically confuse it as some other object of a similar size. As examples, even though very few children identified the pencil sharpener as a pencil sharpener, many children said that it looked like another small object (i.e., binoculars, camera), and two children misidentified the grill as a desk. So, we re-coded the basic level guesses not for correctness but for real-world size. Next the displays were again divided into two groups: (1) displays in which either object's size was correctly guessed at a rate above the median across all items (both items >87.5% correct, 8/20 pairs, M = 97.7% across items), and (2) displays that fell below the median (one or both items < 87.5% correct, 12/20 pairs, M = 80.6% across items). This grouping method was used because size-identification was relatively high overall. As before, we found that pairs of objects whose sizes were poorly identified generated equivalent Size-Stroop effects to pairs with objects whose sizes were well identified (t(18)=-1.42, p = 0.17).

Taken all together, these analyses show that there is remarkable consistency in the specific pairs of big and small objects that generate Size-Stroop effects in adults and children, and that explicit recognition of these objects is not a major mediating factor in the

pairs of objects generated equivalent or larger Size-Stroop effects that well-identified pairs of objects.

Size-Stroop effect for children. These analyses begin to shed light on the mechanisms supporting the Size-Stroop effect in children, indicating that similar mechanisms support real-world object size representations in adults and preschoolers. Broadly, they are consistent with the hypothesis that preschoolers have abstracted the mid-level perceptual features that distinguish big from small objects, as classes, and that preschoolers' visual systems automatically use these features to compute the real-world size of pictured objects.

General Discussion

Across two experiments, we found evidence to suggest that preschoolers automatically activate real-world size information when they see pictured objects. Preschoolers were impaired at making visual size judgments about pictured objects when the relative sizes of the images were incongruent with their relative sizes in the real-world, even though the real-world size was not relevant to the task. This effect was evident in preschoolers' error patterns and reaction times: 3- and 4-year-olds made more errors on incongruent displays, and 4-year-olds also took longer to make visual size judgments on incongruent displays. Further, item-pair analyses showed that the same pairs of big and small objects generated stronger Size-Stroop effects in children and adults, regardless of how well children could identify these depicted objects. Taken together, these results suggest that object size is automatically encoded in preschoolers' object representations and points towards the idea that similar mechanisms may underlie object size representations in both preschoolers and adults.

How do children compute real-world size information?

It could have been the case that preschoolers showed the Size Stroop effect, but that different pairs of big and small objects generated stronger Size Stroop effects in adults and

children. Instead, we found a convergence in Stroop display effects across adults and children. Thus, these results provide *indirect* evidence that children, just like adults, use mid-level perceptual features to directly infer the real-world sizes of objects (Long et al, 2016; Long & Konkle, 2017).

One obvious empirical route for confirming whether preschoolers' visual systems can also use mid-level perceptual features to infer real-world size would be to see if children show the Size-Stroop effect with unrecognizable texforms. Unfortunately—as might be predicted by the fact that children rarely completed these studies with recognizable objects —pilot studies showed that children would not sit through paradigms when stimuli were meaningless blobs.

We see two ways of exploring this hypothesis. The first way would be to specify the mid-level features that are reliable cues to real-world size, and to show that the presence or absence of these features explains the item effects we see both with children and adults. Delineating the key perceptual features that distinguish big objects from small objects, however, is still an area of active research (Long et al., 2016; Long & Konkle, 2017), so pursuing this empirical approach awaits further progress on that front.

The second way would be to follow up the suggestive data from Experiment 3 that children's recognition of objects in terms of basic-level or superordinate-level kinds is irrelevant to the effects of item-pairs. Given that we did not explicitly manipulate how recognizable the items in the displays would be to children, it is possible that a greater disparity in recognition across Stroop displays could reveal a different pattern of results. In other words, the relative contributions of information derived from perceptual features vs. object kind recognition to automatic size processing still remains an open question. Nonetheless, a straightforward prediction of the hypothesis that preschool children's Size-

Stroop effect, like that of adults, can be driven by mid-level perceptual features alone would be that children infer the real-world size of an object they cannot recognize at a basic or superordinate level.

Does real-world size organize neural responses to objects in preschoolers?

Given that preschoolers show the same behavioral signatures of real-world size representation as adults, one possibility is that preschoolers also show the same neural signatures of object size representation. In adults, large swaths of object-selective cortex respond more strongly to pictures of small objects than big objects, and other regions show the opposite preference (Konkle & Oliva, 2012b; Konkle & Caramazza, 2013; Julian, Ryan, & Epstein, 2016); these preferences are stable across changes in the retinal sizes of objects (Konkle & Oliva, 2012b) and are also elicited by unrecognizable versions of big and small objects (Long, Yu, & Konkle, 2017).

Thus, future neuroimaging studies could explore whether a large-scale organization of object-selective cortex by real-world object size is already in place by the preschool years. The few neuroimaging studies in infants and 5–7 year old children suggest that this may be plausible. For example, early regional preferences for faces vs. houses are evident in infants (Deen et al., 2017) but the degree of selectivity of these regions develops through childhood and into adolescence (Golarai, Liberman, Grill-Spector, 2015; Gomez et al., 2018; Natu et al., 2017). One interesting possibility is the broader division of big and small objects emerges early and helps to scaffold further category-selective responses. Consistent with this idea, the general overall similarity structure in neural responses to faces, objects, bodies, and scenes is similar between adults and children 5-7 years of age (Cohen et al., 2016), despite the further refinements to come in adolescence related to

facial recognition and reading/writing abilities (e.g., Carey & Diamond, 1994; James et al., 2017).

Might younger children also show the Size-Stroop effect?

When and how do children begin to automatically process the real-world sizes of pictured objects? One possibility is that younger infants and toddlers may first need to access basic-level representations (e.g., "bottle") before they can access size representations. In doing so over the first few years of life, they may then learn the midlevel features that are implicated in the processing of real-world size. However, another possibility is that mid-level perceptual representations may become linked to real-world size processing relatively early in life. Infants could acquire the perceptual representations that characterize big vs. small objects as classes without the need for basic-level kind representations, possibly as a byproduct of visual and haptic experience with objects of different sizes (Granrud, Haake, & Yonas, 1985). If so, then the perceptual features of unfamiliar objects could already activate real-world size information in young infants.

As this Size-Stroop paradigm was already difficult to run with 3-year-olds, future research will need to develop new methods to examine if and how younger children activate real-world size information when they see pictured objects. An understanding of the mechanisms that lead to adult-like real-world size representations is not only an important question in developmental cognitive neuroscience, but will inform theories of why and how real-world size organizes our cognitive and neural representations of objects in adulthood.

References

- Bova, S. M., Fazzi, E., Giovenzana, A., Montomoli, C., Signorini, S. G., Zoppello, M., & Lanzi, G. (2007). The development of visual object recognition in school-age children. *Developmental neuropsychology*, *31*(1), 79-102.
- Carey, S., & Diamond, R. (1994). Are faces perceived as configurations more by adults than by children?. *Visual cognition*, *1*(2-3), 253-274.
- Cate, A. D., Goodale, M. A., & Köhler, S. (2011). The role of apparent size in building-and object-specific regions of ventral visual cortex. *Brain research*, *1388*, 109-122.
- Chiou, R., & Lamdon Ralph, M. A. (2016). Task-Related Dynamic Division of Labor Between Anterior Temporal and Lateral Occipital Cortices in Representing Object Size. *Journal of Neuroscience*, *36*(17), 4662-4668.
- Cohen, M., Dilks, D., Feather, J., Koldewyn, K., Weigelt, S., & Kanwisher, N. (2016). Common representational structures across the ventral visual pathway of children and adults. *Journal of Vision*, *16*(12), 776-776
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior*, 8(2), 240-247.
- Doherty, M. J., Campbell, N. M., Tsuji, H., & Phillips, W. A. (2010). The Ebbinghaus illusion deceives adults but not young children. *Developmental science*, *13*(5), 714-721.
- Ebeling, K. S., & Gelman, S. A. (1988). Coordination of size standards by young children. *Child Development*, 888-896.
- Frank, M. C., Sugarman, E., Horowitz, A. C., Lewis, M. L., & Yurovsky, D. (2016). Using tablets to collect data from young children. *Journal of Cognition and Development*, *17*(1), 1-17.
- Henik, A., Gliksman, Y., Kallai, A., & Leibovich, T. (2017). Size perception and the foundation of Numerical Processing. *Current Directions in Psychological Science*, *26*(1), 45-51.
- James, K. H. (2017). The Importance of Handwriting Experience on the Development of the Literate Brain. *Current Directions in Psychological Science*, *26*(6), 502-508.
- Jolicoeur, P., Gluck, M. A., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, 16(2), 243–275.
- Julian, J. B., Ryan, J., & Epstein, R. A. (2016). Coding of Object Size and Object Category in Human Visual Cortex. *Cerebral Cortex.* 27(6), 3095-3109.
- Jüttner, M., Wakui, E., Petters, D., & Davidoff, J. (2016). Developmental Commonalities between Object and Face Recognition in Adolescence. *Frontiers in psychology*, *7*, 385.

- Gelman, S. A., & Ebeling, K. S. (1989). Children's use of nonegocentric standards in judgments of functional size. *Child Development*, 920-932.
- Gliksman, S. I., Leibovich, T., Melman, Y., & Henik, A. (2016). Automaticity of Conceptual Magnitude. *Scientific Reports*, 6.
- Gomez, J., Natu, V., Jeska, B., Barnett, M., & Grill-Spector, K. (2018). Development differentially sculpts receptive fields across early and high-level human visual cortex. *Nature communications*, 9(1), 788.
- Golarai, G., Liberman, A., & Grill-Spector, K. (2015). Experience Shapes the Development of Neural Substrates of Face Processing in Human Ventral Temporal Cortex. *Cerebral Cortex*, bhv314.
- Gori, M., Del Viva, M., Sandini, G., & Burr, D. C. (2008). Young children do not integrate visual and haptic form information. *Current Biology*, *18*(9), 694-698.
- Granrud, C. E., Haake, R. J., & Yonas, A. (1985). Infants' sensitivity to familiar size: The effect of memory on spatial perception. *Perception & psychophysics*, *37*(5), 459-466.
- Khaligh-Razavi, S. M., Cichy, R. M., Pantazis, D., & Oliva, A. (2018). Tracking the Spatiotemporal Neural Dynamics of Real-world Object Size and Animacy in the Human Brain. *Journal of cognitive neuroscience*, 1-18.
- Konkle, T., & Oliva, A. (2012a). A familiar-size Stroop effect: real-world size is an automatic property of object representation. *Journal of Experimental Psychology: Human Perception and Performance, 38*(3), 561.
- Konkle, T., & Oliva, A. (2012b). A real-world size organization of object responses in occipitotemporal cortex. *Neuron*, 74(6), 1114-1124.
- Konkle, T., & Caramazza, A. (2013). Tripartite organization of the ventral stream by animacy and object size. *The Journal of Neuroscience*, *33*(25), 10235-10242.
- Long, B., Konkle, T., Cohen, M. A., & Alvarez, G. A. (2016). Mid-level perceptual features distinguish objects of different real-world sizes. *Journal of Experimental Psychology: General*, *145*(1), 95.
- Long, B. & Konkle, T. (2017). A familiar Size-Stroop effect in the absence of basic-level recognition. *Cognition*. 168(1), 234-242
- Long, B., Moher, M., Carey, S., & Konkle, T. (under review). Animacy and real-world size are reflected in perceptual similarity computations by the preschool years.
- Long, B., Yu, C.P, & Konkle, T. (in press). A mid-level organization of the ventral stream. bioRxiv. https://doi.org/10.1101/213934
- Natu, V. S., Barnett, M. A., Hartley, J., Gomez, J., Stigliani, A., & Grill-Spector, K. (2016).

Development of neural sensitivity to face identity correlates with perceptual discriminability. *Journal of Neuroscience*, *36*(42), 10893-10907.

- Nishimura, M., Scherf, S., & Behrmann, M. (2009). Development of object recognition in humans. *F1000 biology reports*, *1*.
- Paivio, A. (1975). Perceptual comparisons through the mind's eye. *Memory & Cognition*, *3*(6), 635-647.
- Sellaro, R., Treccani, B., Job, R., & Cubelli, R. (2015). Spatial coding of object typical size: evidence for a SNARC-like effect. *Psychological research*, *79*(6), 950-962.

Whelan, R. (2008). Effective analysis of reaction time data. *The Psychological Record*.

Yonas, A., Pettersen, L., & Granrud, C. E. (1982). Infants' sensitivity to familiar size as information for distance. *Child Development*, 1285-1290.

Appendix



Supplementary Figure 1. Examples of recognizable images and their corresponding texforms, for a group of three big objects (left) and three small objects (right) (Long et al., 2016; Freeman & Simoncelli, 2011).



Supplementary Figure 2. Size-Stroop effects for individual children are shown as a function of children's average age in months (A) and as a function of children's average speed on the task across both conditions (B). The absolute difference between incongruent and congruent conditions (i.e., the absolute value of the Size-Stroop effect) is plotted for each participant as a function of their average speed on the task (C).

A. Image shown



B. 4-year-olds responses

BBQ	Camping food cookingtable
Desk	Machine
Grill	Grill
(no response)	Grill
(no response)	Kitchen oven
Kitchen Table	Something
TV	Grill
Balcony/Cook	Grill
Table	Grill
(no response)	(no response)
Desk	Grill
Cook Thing	A grill

Α.	Image	shown
----	-------	-------



B. 4-year-olds responses

(no response)	Binoculars
Music Box	Pencil Sharpener
(no response)	(no response)
Spy Graph	Radio
(no response)	I don't knownever seen it before
Remote Control	Something with nails in it
Camera	Remote control
Tool Thingy	What you put on your eyes to see far away
Drawer	Binoculars
Binoculars	it looks like a plug
Tool	Binoculars
Instrument	Binoculars

Supplemental Figure 3. Twenty-four 4-year-olds were asked "what does this look like?" about the depicted object shown in (A). Their responses are shown in (B). Responses that were counted as correct recognitions are bolded. Note that responses were coded liberally; for example, "balcony/cook" was accepted as a correct answer for the grill (A, left panel).