Pictorial Depth Cues in Young Children's Drawings of Layouts and Objects

Théo Morfoisse, Todd Gureckis, and Moira R. Dillon
{theo.morfoisse, todd.gurekcis, moira.dillon}@nyu.edu
Department of Psychology, New York University
6 Washington Place, NY 10003, USA

Abstract

Humans have been faced with the challenges of pictorial production since at least the Paleolithic. Curiously, while the capacity to *navigate* layouts and *recognize* objects in everyday life comes almost effortlessly, inherited from our evolutionary past, the capacity to *draw* layouts and objects is more effortful, often needing time to improve over the course of an individual's development and with the technological innovations acquired through culture. The present study examines whether young children might nevertheless rely on phylogenetically ancient spatial capacities for navigation and object recognition when creating uniquely human pictorial art. We apply a novel digital coding technique to a publicly available dataset of young children's drawings of layouts and objects to explore children's use of classic pictorial depth cues including size, position, and overlap. To convey pictorial depth, children appear to adopt several cues, without a preference among them, younger than had been suggested by previous studies that used other, less rich, analytic techniques. Moreover, children use more cues to pictorial depth in drawings of layouts versus objects. Children's creation of uniquely human pictorial symbols may thus reflect their heightened use of depth for navigating layouts compared to recognizing objects, both cognitive capacities that humans share with other animals.

Keywords: drawings; child development; layouts, objects; navigation; depth

Introduction

Pictorial depth cues in human art have varied across history and culture. For example, while Chinese landscapes from the third century CE incorporated few cues to pictorial depth, artists from the twelfth century in the Song dynasty could rely on centuries of accumulated ink-wash and brush stroke techniques to convey atmospheric depth (Ping, 2015). In the West, the incorporation of vanishing points to create perspectival depth relied, at least in part, on technologies like lenses and mirrors and allowed for more accurate renderings of the depth information describing both layouts and objects (Falco, 2016). While history has revealed great variety and innovation in the way humans *depict* the world, science has revealed striking similarities in the way humans and even other animals *navigate* the world. Moreover, since at least the time of Piaget, psychologists and cognitive scientists have examined the art of young children for insight into the cognition driving this uniquely human act (e.g., Piaget & Inhelder, 1956). Children's art may reveal what challenges of artistic representation can be overcome through universal development and experience and what challenges persist regardless of such universal experience.

For example, young children may intuitively rely on certain techniques to convey pictorial depth in their drawings, and, although the literature is mixed, these techniques may change through development. Psychologists have explored children's use of pictorial depth cues like the relative size of objects, objects' positions on the picture plane, and whether or not objects overlap. While even five-year-old children can capture the relative, ordinal sizes of two objects in their drawings (Braine, Schauble, Kugelmass, & Winter, 1993; Silk & Thomas, 1988), some studies suggest that children do not consistently use size to indicate depth in their drawings (i.e., making an object in the background smaller than an object in the foreground) until around age nine years (Cox & Perara, 2001). Indicating depth using the relative positions of objects or whether they overlap may also undergo developmental change: When asked to draw two apples, one behind the other, five-year-old children tend to draw the apples side-by-side, eight-year-old children draw one higher on the page than the other, and ten-year-old children draw them overlapping (Freeman, Eiser, & Sayers, 1977).

Curiously, most studies probing children's use of depth cues in drawings ask children to draw only objects, and those that do probe use of depth cues to depict both layouts and objects assume that depth is treated the same way across these two spatial contexts (e.g. Cox & Littleton, 1995). Moreover, studies focused on children's use of depth cues when drawing layouts nevertheless ask children to draw from either photographs or scale models, not large, 3D environments that are actually navigable (Ebersbach, Stiehler, & Asmus, 2011; Lange-Küttner, 2014). These studies thus fail to capture critical contrasts between these spatial categories: Layouts are large spaces that can be walked through; objects are composed of parts that are small and manipulable (see Epstein and Kanwisher, 1998; Spelke & Lee, 2012). As a result, these studies fall short in addressing whether human and other animal's intuitive and differential use of depth information for navigating layouts and recognizing objects is expressed in uniquely human drawings.

In contrast to prior work, recent work by Dillon (2019) presented young children with a large navigable "fort" or a small toy model of that fort (*Fig. 1*) and asked children to draw "exactly what they see." Dillon (2019) found that children often omitted the fort's walls but included the corresponding object-part information for the toy. Spatial category elicited children's different drawing behaviors since the fort and toy corresponded exactly on shape and configuration. Dillon (2019) suggests that children omit the fort's walls because humans and other animals use such extended layout features in the environment automatically

during navigation (Doeller & Burgess, 2008), but drawings and other symbolic media, like language, prioritize elements in a scene that are explicitly, not automatically, attended to (Gentner, 1982; Scholl, 2001). Such differences in drawing layouts versus objects, rooted in the everyday navigation and object recognition shared by humans with other animals, may also be present in children's use of specific spatial cues, like depth cues, in drawings. In particular, when children do make layout information explicit by including it in their drawings, they may adopt more cues to pictorial depth given depth's relevance to navigation.

The present study explores this possibility. It develops a novel digital coding technique that permits precise quantitative analyses of children's drawings of layouts and objects and applies it to a publicly available dataset of children's drawings (https://osf.io/5wng2/). Such quantitative tools allow us to explore whether children systematically use size, position, and overlap in their drawings in a way that reflects their use of depth information for navigating layouts versus recognizing objects.

Methods

The present study analyses the drawing set by Dillon (2019) available on the Open Science Framework at: https://osf.io/5wng2/. Full descriptions of that study's methods are also available.

Participants

The dataset includes two drawings each from thirty-two four-year-old children (*M*age = 4.50, range = 4.03 - 4.99; 15 females) assigned to draw a "fort" (*Fig. 1*) and two drawings each from a different group of 32 four-year-old children (*M*age = 4.49, range = 4.02 - 4.98; 21 females) assigned to draw a "toy."

Drawing Stimuli and Instructions

In the original study, children were presented with four configurations of either a colorful fort or toy. Two configurations included 3 walls and 3 objects, and two configurations included 4 walls and 4 objects. The present study just considers the 3-wall/3-object configurations (*Fig. 1*) because their spatial information was more regular. In the 3-wall/3-object configurations, the fort presented two rectangular side walls (5'6" x 7') and one rectangular back wall (5'6" x 5'3"), with one rectangular object (2' x 1'6") in front of each wall. One of these configurations also included circular decals placed in the center of each wall (diameter = 1'9") and object (diameter = 6"). During drawing, children sat 6'' in from the fort's opening, and the back wall of the fort subtended 44.69 dva in the vertical direction and 44.12 dva in the horizontal direction.

The toy presented the same configurations as the fort, but at 1/20 the size. During drawing, children sat on a chair in front of the toy, and the toy was covered as children walked into and out of the room so they could never see it from above. During drawing, children viewed the toy at eye-level, 12" away, and the back wall of the toy subtended 11.46 dva in the vertical direction and 10.89 dva in the horizontal direction.

In both conditions, children were instructed: "draw exactly what you see." They completed one practice drawing, where they copied a 2D picture, followed by the test drawings, one for each configuration, in a semi-random order. After children indicated that a drawing was complete, the experimenter asked them to point to each individual element in their drawing and to identify it by touching what it referred to in the fort or on the toy. The child thus provided an identifying label for each element they drew.

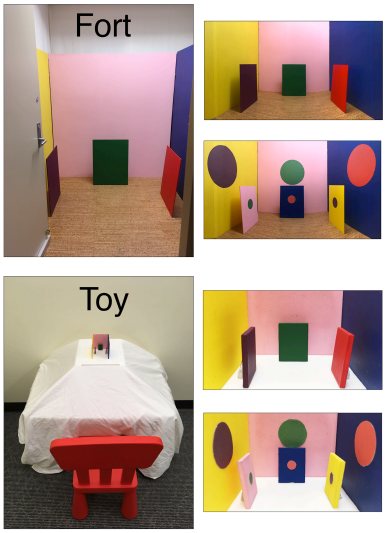

Fig. 1. Photographs of the context and configurations in the fort and toy conditions. Adapted from Dillon (2019).

Coding and Analyses

The novel digital coding technique developed for the present study offers significant advantages compared to previous coding techniques, especially in allowing us to evaluate the precise spatial properties of the drawings. For example, in their evaluation of children's drawings of two objects that were described as near or far or in front or behind, Braine et al. (1993) judged differences in area and horizontal or vertical position by eye, or in the cases where area differences were small, by measuring at least a 10% difference in the length of a drawn object's principal axis. Their coding, moreover, used a simple categorical ordering of the two elements' sizes, as opposed to a continuous size metric. Studies evaluating whether drawn elements overlap (e.g., Braine et al., 1993; Freeman et al., 1977) have used a similar categorical judgment. Some prior studies have evaluated the height of drawn elements by measuring their principal axis with a ruler (e.g., Lange- Küttner, 2009; 2014) or the area of drawn elements by overlaying drawings with a transparent grid (e.g. Silk & Thomas, 1988). These coding strategies nevertheless fall short in providing any precise area measurements. Finally, Konkle & Oliva (2011) used an automatic digital coding technique that identified the bounding box of a drawn element, and size was evaluated by the length of the bounding

box's diagonal. Nevertheless, the drawings in this study were more amenable to automatic coding than the drawings in the present study since they were made by adults and included only one object per page. Their metric, moreover, still did not provide any precise area measurements.

Our coding technique thus offered significant advances in precision and robustness compared to previous techniques. Research assistants, unaware of any planned analyses, carefully traced each element in the original scanned drawings using the Pen Tool in Adobe Illustrator. Elements were traced either as lines or polygons, depending on whether they were open or closed. We used the child-generated labels from the original study to identify and label each traced element. SVG files, which contained the label of each element and the coordinates of its boundaries, were exported from Illustrator for each drawing. Such rich information allowed for the automatic calculation of many metrics describing the size, position, and overlap of the drawn elements. Analyses were conducted in Matlab (R2018 b) and R (Version 1.2). Eight drawings from each condition were recoded by a second coder, and reliability, measured by comparing the absolute area of the outlined elements from the first and second coders collapsed across the fort and toy conditions, was very high (ICC(1,1) = .98, 95% CI [.97, .99]).

Results

Using our digital coding, we evaluate how size, position, and overlap might serve as depth cues in young children's drawings of layouts and objects.

Size

We first tested whether children captured size information accurately in their drawings by evaluating the absolute and relative area occupied by the walls and objects in each drawing across conditions. We considered drawings in which there was at least one wall drawn (fort: N = 40; toy: N = 61).

A mixed-model linear regression with condition (fort or toy) and spatial element (wall or object) included as predictor variables and participant included as a random-effects intercept revealed that children drew elements absolutely bigger in the fort versus toy condition ($p = .042$, $\beta = 0.23$) and drew walls bigger than objects ($p = .004$, $\beta = 0.28$; *Figs. 2-3*). A significant interaction revealed that the difference in area between the walls and objects was bigger in the fort versus toy condition ($p = .020$, $\beta = 0.36$).

To test whether children drew the walls and objects according to their real-world proportions, we computed the relative area of the objects compared to the walls in the real world and in children's drawings. While objects were 8% of the area of the walls in the real world, children drew the objects at 36% of the area of the walls in the fort condition and 39% in the toy condition. Given the average area that the walls were drawn in the fort (35861 px^2) and toy (14485 px^2), the objects should have had an average area of 3190 px^2 and 1218 px^2. The actual average sizes that the objects were drawn (fort: 13041 px^2; toy: 5626 px^2) were significantly different from these values (intercept-only mixed-model linear regression, fort: $t(20) = 2.92$, $p = .007$; toy: $t(30) = 2.58$, $p = .015$).

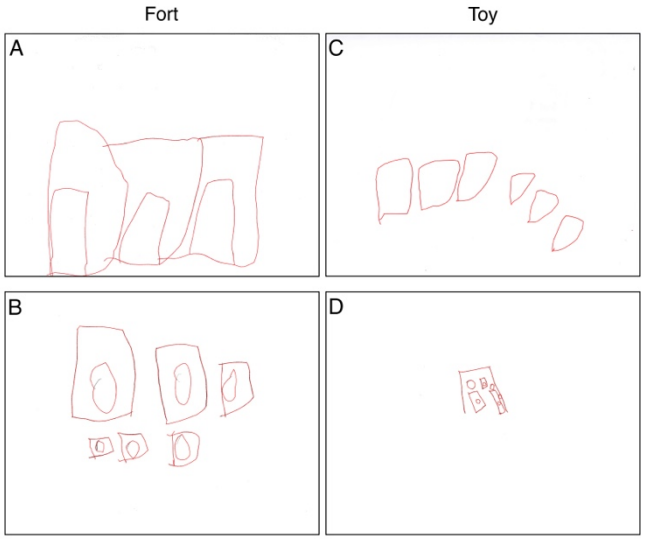

*Fig. 2. Example drawings from different children in the fort (**A**-**B**) and toy (**C**-**D**) conditions. Drawings **A** and **C** are of the 3-wall/3-object configurations with no decals, and drawings **B** and **D** are of the 3-wall/3-object configurations with decals (see Fig. 1). Children in the fort condition tended to draw bigger elements, and the objects sometimes overlapped the walls. Children in the toy condition, in contrast, tended to draw smaller elements, and the objects rarely overlapped the walls. Children in both conditions used the elements' horizontal (e.g., **C**) and vertical (e.g., **B**) positions as cues to pictorial depth, although children in the fort condition used both of these positional relations significantly more.*

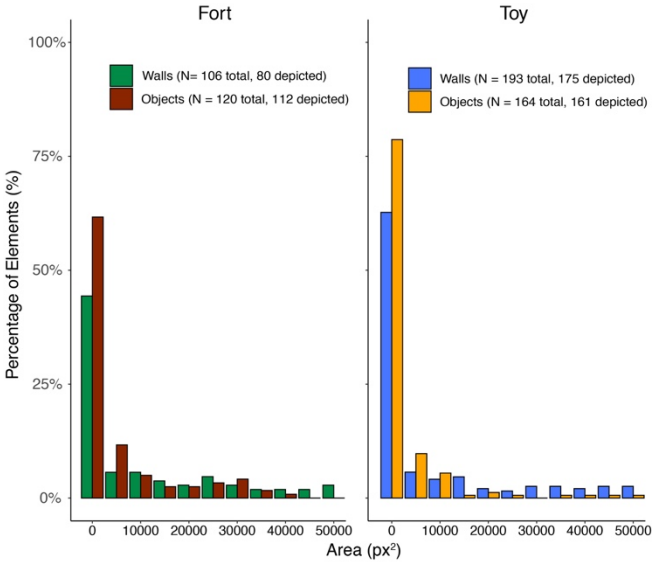

Fig. 3. Distribution of the percentage of elements at different areas. Because these distributions have a very long tail, only 75% of the data are depicted here. The absolute area of drawing surface is 484704 px^2 (792 px x 612 px). Bin size = 5000 px^2.

Position

In prior research, young children indicated that one object was behind another by placing the two objects side-by-side in their drawing (e.g., Braine et al., 1993; Freeman et al., 1977; see Fig. 2, toy condition). Considering drawings with at least one wall and one object and with < 10% overlap between objects and walls (fort: N = 21; toy: N = 39; see **Results**, **Overlap**), we calculated the horizontal position of each element in each drawing using the x-coordinate of the element's center of mass. 19.05% of drawings in the fort condition and 7.69% of drawings in the toy condition depicted all walls on one side of all objects. A mixed-model binomial logistic regression including condition as a predictor variable and participant as a random-effects intercept revealed a significant difference between conditions (p = .020, β = 12.23).

Older children in prior work, in contrast, relied more on vertical instead of horizontal position to indicate depth, drawing more distant elements higher on the page than more proximal elements. Considering drawings with at least one wall and one object and with < 10% overlap between objects and walls (fort: N = 21; toy: N = 39; see **Results**, **Overlap**), we calculated the vertical position of each element in the drawing using the y-coordinate of the element's center of mass. 19.05% of drawings in the fort condition and 10.26% of drawings in the toy condition depicted all walls higher than all objects. A mixed-model binomial logistic regression including condition as a predictor variable and participant as a random-effects intercept revealed a significant difference between conditions (p = .029, β = 9.82).

Overlap

In prior work, mostly older children (age around ten years) overlapped elements in their drawings as a cue to pictorial depth. We tested whether children in the present study drew objects that overlapped walls by computing the area of each object in each drawing that overlapped any wall. We then converted that area to a percentage of the object's total area. If an object was completely overlapping a wall, this value would be 100% (e.g., *Fig. 2**A***), and if an object overlapped no walls, this value would be 0% (e.g., *Fig. 2**C***). We considered all drawings with at least one wall and one object (fort: N = 37; toy: N = 57).

Fig. 4 displays the distribution of the percentage of object overlap in both conditions. Strikingly, the distribution is largely bimodal, with most objects overlapping walls very little or not at all (< 10%), and some objects completely or almost completely overlapping walls (> 90%). In the fort, 29% of objects fell into this latter category (completely or almost completely overlapping) compared to only 17% of objects in the toy condition. Consistent with this difference in percentages, a mixed-model linear regression with the percentage of object overlapping as the dependent variable, condition and absolute object area as predictor variables, and participant as a random-effects intercept revealed greater overlapping in the fort versus toy conditions (p < .001, β = 0.21) after controlling for object area, which was also significant (p = .018, β = 0.17).

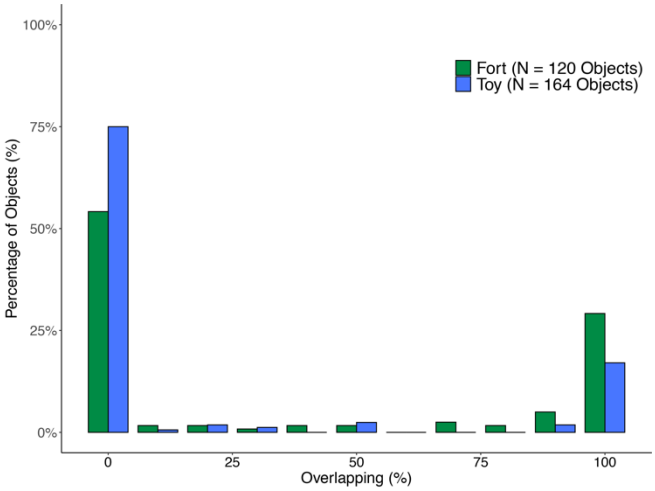

Fig. 4. Distribution of the percentage of object overlap. 100% indicates objects totally overlapping walls. Bin size = 10%.

Discussion

A precise quantitative analysis of the spatial information in young children's drawings of layouts and objects revealed a more precocious and more nuanced picture of young children's use of pictorial depth cues than had been suggested by prior research. In particular, while prior studies may have suggested a developmental trajectory in the use of pictorial depth cues, our analyses revealed that even young children use a variety of pictorial depth cues in their drawings. Our analysis also revealed striking differences in the ways in which children use these depth cues in drawings of layouts versus objects, two spatial contexts that are cognitively and neurally dissociated for both humans and other animals.

Children were successful in capturing the real-world, ordinal size relations of the layout and object elements that they drew, drawing walls bigger than objects and forts bigger than toys. Nevertheless, they failed to capture the precise metric size relations that described the walls and objects and drew objects bigger, relative to walls, compared to their real-world proportions. Moreover, children's use of size as a pictorial cue did not differ between the fort and toy conditions.

Children also used the horizontal and vertical positions of walls and objects in their drawings in consistent ways. Children sometimes drew walls and objects side-by-side and sometimes drew walls above objects. Both of these tendencies were stronger in the fort versus toy condition.

Finally, children employed overlap in their drawings, drawing objects overlapping walls instead of free-standing. This tendency also occurred more often in the fort versus toy condition.

Below we first discuss the methodological advances of the current study. Then, we discuss how our findings for size, position, and overlap might serve as cues to pictorial depth

and might reflect children's reliance on depth information for everyday navigation.

Methodological Advances

Our coding technique offered significant advances in precision and robustness compared to previous techniques used to analyze drawings. Nevertheless, coding was laborious, time-consuming, and completable only by trained human coders. Moving forward, automizing at least some of the coding process would be preferable, though how some of the available techniques and technologies could be adopted is uncertain. Experimenters might consider having children draw with digital tools instead of a paper and pencil, but for children as young as the ones in the current study such tools may be impractical: Young children often place their non-drawing hand on the drawing surface, move their fingers along the drawing surface, and get distracted by standard tablet and stylus buttons — actions that might affect digital drawing. Moreover, while young children often have at least some experience with a paper and pencil, few have experience with the particular digital tablet and stylus used in a given study, and such novelty makes it harder for young children to engage with the experimental task. Digital notepads, like the Bamboo series from Wacom, where children can draw with a pen on paper but where their marks are automatically converted to a series of vector graphics, may be a good solution for future research.

Some prior work investigating children's drawings has successfully used methods from computer vision to automatically recognize and evaluate children's drawings of objects (e.g., Long, Fan, Chai, & Frank, 2019). The automaticity of such methods is hugely advantageous, but such tools are currently ill-equipped for datasets like the one used in the present study. First, these tools are trained on datasets of real-world objects (e.g., those from Imagenet), which have distinctive shapes and features. In the present study, the fort and toy were composed only of rectangles whose configurations were matched across spatial contexts. As a result, not only was there no distinctive shape information to differentiate between contexts and elements, but also there was no way to use such shape features to label the different elements. Second, computational methods do best with large datasets (e.g., Long et al., 2019, evaluated thousands of drawings), and the present dataset had relatively few drawings. Despite these challenges, the present digital coding technique, with its capacity to capture and reveal precise and consistent quantitative differences in the way children depict layouts versus objects, may inform future computational models aimed at identification and classification of spatial information, even in children's drawings of such simple configurations.

Size

Four-year-old children in both the fort and toy conditions drew walls bigger than objects, consistent with prior literature suggesting that young children correctly capture the ordinal real-world sizes of objects in their drawings (Silk & Thomas, 1988). More surprisingly, children in the fort condition made drawings with elements that were absolutely bigger than the elements drawn in the toy condition even though different groups of children participated in these different conditions. These results are consistent with the drawings of adults who depicted objects of different real-world sizes (Konkle & Oliva, 2011) and suggest that the drawing space itself may serve as a consistent scaling factor for real-world size. Future studies might explore whether the actual, real-world size information that differentiated between the fort and toy is necessary for this result or if referential context alone, e.g., referring to the very same spatial arrangement as a navigable space or a manipulable object, may also affect the size that children draw layout and object elements (see DeLoache, 2004). Young children's interpretation of drawings is affected by such referential context: When young children are shown the very same drawing, described as representing either a navigable layout or a manipulable object, they use different geometric information to find a hidden location in the layout or on the object using the drawing (Dillon, Huang, & Spelke, 2013; Dillon & Spelke, 2017).

Children effectively captured information about the real-world size of layouts and objects in their drawings, but they may have also used size as a cue to pictorial depth. Consistent with this possibility, children depicted objects proportionally bigger than they should have, given the size that they drew walls. Children as young as age five years sometimes draw elements in the foreground bigger than elements in the background, although previous research has only tested this capacity in children's drawings of objects, not of objects and layouts (Cox, 2005; Cox & Perara, 2001; Silk & Thomas, 1988). Because children drew objects proportionally bigger in both the fort and toy conditions, however, it is likely that their use of size information applies generally across different spatial contexts. Dillon (2019) found that while the inclusion of wall or object information varied greatly according to spatial context (with few walls being drawn in the fort condition, but the corresponding object-parts consistently being drawn in the toy condition), the dimensionality with which walls and objects were drawn (i.e., as single lines or closed, frontoparallel figures) varied based on what was in the foreground or background, not on spatial category. The present results suggest that drawn element size may function similarly, with elements in the foreground being drawn proportionally bigger, regardless of spatial context.

Although we suggest that children in the present study used size as a cue to pictorial depth, one possible alternative is that children simply drew the foreground elements first, leaving less room for appropriately scaled walls (Freeman, 1980; Lange-Küttner, 1997; 2004; 2009). Future studies using digital tools, capable of recording the order with which children draw different elements could better address this suggestion.

Position

Children in both the fort and toy conditions used the horizontal and vertical positions of walls and objects in their

drawings in consistent ways. While according to the previous literature the exact age at which children adopt these positional relations as cues to pictorial depth varies somewhat, children in the present study are among the youngest who have shown use of these cues (Braine et al., 1993). Moreover, while some studies suggest that young children preferentially use horizonal position versus vertical position as cues to pictorial depth (e.g., Freeman et al., 1977), children in the present study appear to use vertical position just as frequently as horizontal position.

Four differences between the present study and prior work are worth noting when interpreting these results. First, the analysis in the present study considered drawings with little-to-no element overlap and used the x- and y-coordinates of an element's center of mass to determine its position. Most prior studies used simple visual judgements of position made by human coders. Although we believe our specifications are well-matched to human judgments (and visual inspection of the classified drawings corroborate the findings of our analysis), many other metrics could also have been chosen and could have affected the frequency with which positional cues were observed. Future analyses might thus focus on evaluating element position using different metrics (especially since many are derivable from the rich information provided by the existing coding).

Second, our analysis allowed for the very same drawing to count as having both horizontal and vertical cues to pictorial depth. Prior studies either used an exclusive binary classification (e.g., Freeman et al., 1977) or used an exclusive tertiary classification, including horizontal, vertical, and diagonal position categories, with the diagonal category being judged as some combination of the horizontal and vertical category (e.g., Braine et al., 1993).

Third, children in the present study were asked to draw elements with very simple shapes, only rectangles and circles. Children in prior studies, in contrast, were asked to draw more complicated shapes, like apples, dogs, people, and houses. Fewer demands to draw complicated shape information paired with fewer shape-based properties that could distinguish among the drawn elements could have led children in the present study to include more positional information in their drawings.

Fourth, prior studies examining position as a cue to pictorial depth in children's drawings have focused on children's drawings of objects, not of layouts and objects. Children in the present study used both horizontal and vertical position more in the fort versus toy condition. Asking children to draw a navigable layout may thus elicit more depth cues in drawings. Humans and other animals alike prioritize egocentric depth or distance information to the boundaries of the extended layout during everyday navigation (see, e.g., Cheng & Newcombe, 2005; Epstein, Patai, Julian, & Spiers, 2017; Spelke & Lee, 2012 for reviews). And so, when children make layout information explicit in their drawings by drawing the walls of the fort, they may include more cues to pictorial depth given depth's relevance to navigating layouts.

Finally, Cox (2005) suggests that even if children use vertical position in their drawings, it does not mean that this positional information is a cue to pictorial depth. Rather, children may just be drawing closer elements in the real world closer to them on the page. The present study offers counterevidence to this proposal. The fort and toy conditions were matched on their configurations, including on the relations between their foreground and background elements. If Cox's (2005) suggestion were true, then there should have been no difference in children's use of vertical position in the fort and toy conditions.

Overlap

Young children in the present study also adopted what is generally agreed upon in the prior literature as a more advanced cue to pictorial depth: overlap (e.g., Braine et al., 1993; Cox, 2005; Freeman et al. 1977). Overlap was also included more often in the fort versus toy condition. As suggested above, the rich coding and metric of evaluation, the use of simple shapes, and the inclusion of both layout and object information may have led to the result that children included overlap as a cue to pictorial depth more in the present study compared to prior studies. Moreover, when children make layout information explicit in their drawings by drawing the walls of the fort, they may include more overlap as a cue to pictorial depth given depth's relevance to navigating layouts.

Conclusion

The present study develops a novel digital annotation technique that permits a precise and rich quantitative analysis of children's drawings. When applied to a dataset of children's drawings of layouts and objects, our analysis suggests that children incorporate spatial information about size, position, and overlap, all cues to pictorial depth, at least as young as if not younger than had been suggested by previous studies. Moreover, children use these cues with no clear preference or progression.

Our analysis also revealed some similarities but also striking differences in the ways that children convey pictorial depth in drawings of layouts versus objects. For example, Dillon (2019) showed that children use dimensionality similarly when drawing layouts and objects; i.e., children depict elements in the background more often as lines versus frontoparallel figures compared to elements in the foreground, regardless of the spatial context. Consistent with this finding, the present analysis revealed that children used size similarly regardless of spatial category; i.e., they depicted foreground elements proportionally bigger relative to background elements given the elements' real-world sizes.

Dillon (2019)'s main finding, however, focused on the differences between children's depictions of different spatial contexts; in particular, children's exclusion of layout versus object information in their drawings. Dillon (2019) suggests that children omit the fort's walls because everyday navigation relies automatically on the layout distance information defined by the walls, but drawings communicate

what is explicitly, not automatically, attended to. The present findings are consistent with these results. In particular, when children *do* make layout information explicit by including it in their drawings, they *also* make depth information explicit by including cues in their drawings like position and overlap. While history reveals innovative variety in the ways humans depict the world, cognitive science may uncover phylogenetically ancient and early developing cognition that both supports and constrains these depictions.

Acknowledgements

This work was supported by a National Science Foundation CAREER Award (DRL-1845924 to M.R.D.) and by a Jacobs Foundation Early Career Fellowship (to M.R.D.). Thanks to P. Chan, A. Bachrach, and I. Cekici for assistance with coding, M. Thornton for assistance with data analysis, and B. Reilly for suggestions on the manuscript.

References

Braine, L. G., Schauble, L., Kugelmass, S., & Winter, A. (1993). Representation of depth by children: Spatial strategies and lateral biases. *Developmental Psychology*, *29*(3), 466-479.

Cheng, K., & Newcombe, N. S. (2005). Is there a geometric module for spatial orientation? squaring theory and evidence. *Psychonomic Bulletin & Review*, *12*(1), 1–23.

Cox, M., & Littleton, K. (1995). Childrens Use of Converging Obliques in their Perspective Drawings. *Educational Psychology*, *15*(2), 127–139.

Cox, M. V., & Perara, J. (2001). Children's use of the height and size cues to depict a projective depth relationship in their pictures. *Psychologia: An International Journal of Psychology in the Orient, 44*(2), 99–110.

Cox, M. (2005). *The pictorial world of the child*. Cambridge: Cambridge University Press.

DeLoache, J. S. (2004). Becoming symbol-minded. *Trends in Cognitive Sciences*, *8*(2), 66–70.

Dillon, M. R. (2019, October 23). Rooms without walls: Young children draw objects but not layouts. https://doi.org/10.31219/osf.io/zqc6k

Dillon, M. R., Huang, Y., & Spelke, E. S. (2013). Core foundations of abstract geometry. *Proceedings of the National Academy of Sciences*, *110*(35), 14191-14195.

Dillon, M. R., & Spelke, E. S. (2017). Young children's use of surface and object information in drawings of everyday scenes. *Child Development, 88*(5), 1701-1715.

Doeller, C. F., & Burgess, N. (2008). Distinct error-correcting and incidental learning of location relative to landmarks and boundaries. *Proceedings of the National Academy of Sciences*, *105*(15), 5909–5914.

Ebersbach, M., Stiehler, S., & Asmus, P. (2011). On the relationship between children's perspective taking in complex scenes and their spatial drawing ability. *British Journal of Developmental Psychology*, *29*(3), 455–474.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598–601.

Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: spatial navigation and beyond. *Nature Neuroscience*, *20*(11), 1504–1513.

Falco, C. M. (2016). Optics and Renaissance Art. *Optics in Our Time*, 265–283.

Freeman, N. H. (1980). *Strategies for representation in young children: Analysis of spatial skills and drawing processes*. UK: Academic Press.

Freeman, N., Eiser, C., & Sayers, J. (1977). Children's strategies in producing three-dimensional relationships on a two-dimensional surface. *Journal of Experimental Child Psychology*, *23*(2), 305–314.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. A. Kuczaj (Ed.), *Language development*: *Language, thought, and culture* (Vol. 2, pp. 301-334). Hillsdale, NJ: Erlbaum.

Konkle, T., & Oliva, A. (2011). Canonical visual size for real-world objects. *Journal of Experimental Psychology : Human Perception and Performance, 37*(1), 23–37.

Lange-Küttner, C. (1997). Development of size modification of human figure drawings in spatial axes systems of varying complexity. *Journal of Experimental Child Psychology*, *66*(2), 264–278.

Lange-Küttner, C. (2009). Habitual size and projective size: The logic of spatial systems in children's drawings. *Developmental Psychology*, *45*(4), 913–927.

Lange-Küttner, C. (2014). Do drawing stages really exist? Children's early mapping of perspective. *Psychology of Aesthetics, Creativity, and the Arts*, *8*(2), 168–182.

Long, B., Fan, J., Chai, Z., & Frank, M. C. (2019). Developmental changes in the ability to draw distinctive features of object categories. *Proceedings of the 41st Annual Conference of the Cognitive Science Society.*

Piaget, J., & Inhelder, B. (1956). The child's conception of space. London, UK: Routledge & Kegan Paul.

Ping, F. (2015). *The efficacious landscape on the authorities of painting at the Northern Song court*. Cambridge, MA: Harvard Univ. Asia Center.

Scholl, B. J. (2001). Objects and attention: the state of the art. *Cognition*, *80*(1-2), 1–46.

Silk, A. M. J., & Thomas, G. V. (1988). The development of size scaling in children's figure drawings. *British Journal of Developmental Psychology*, *6*(3), 285–299.

Spelke, E. S., & Lee, S. A. (2012). Core systems of geometry in animal minds. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1603), 2784–2793.