Spatial language and intuitive physics in children and adults: It's not so simple

Karima Elgamal (karima02@bu.edu)
Department of Psychological & Brain Sciences,
64 Cummington Mall, Boston MA 02215 USA

Rennie Pasquinelli (rennpasq@gmail.com)
Department of Cognitive Science,
237 Krieger Hall 3400 N. Charles St, Baltimore MD 21218 USA

Laura Lakusta (lakustal@montclair.edu)
Department of Psychology,
Dickson Hall 247 1 Normal Ave, Montclair NJ 07043 USA

Barbara Landau (landau@jhu.edu)
Department of Cognitive Science,
237 Krieger Hall 3400 N. Charles St, Baltimore MD 21218 USA

Abstract

Do simple spatial terms such as *in* or *on* map directly to intuitive physical judgements about spatial relationships between objects that underlie these terms' meaning? We explored this question in the domain of physical support. Adults (N=120) and 4-year-old children (N=42) were shown videos in which a puppet placed an L-shaped object in contact with a table at locations that varied in whether the object was supported or not. Half of the participants were asked for linguistic judgments ("Is X on Y?") and half were asked for intuitive physics judgments ("Will X fall if (agent) lets go?"). Results revealed that linguistic judgments were largely categorical, with child and adult participants labeling objects as *on* even when the object was not truly supported. In contrast, intuitive physics judgments aligned closely with the object's actual possibility of true support. However, responses also varied by the orientation of the L-shaped object, with *on* applying categorically to a regularly oriented L, but in a more graded fashion for a mirror image oriented L. Our findings suggest that the mapping between the simple spatial term *on* and physical reasoning systems are not completely coupled, and that the ways in which language draws on intuitive physical reasoning is complex.

Keywords: spatial language; intuitive physics; support relations; spatial reasoning

Introduction

A prominent issue in our understanding of spatial terms in adults and their acquisition in children concerns the relationship between the meanings of spatial terms and their non-linguistic foundation. Does non-linguistic knowledge serve as the foundation for learning individual spatial terms such as *in* and *on*? These terms are acquired within the first few years, and by hypothesis, young children might acquire the term *on* when they have achieved an intuitive pre-linguistic notion of support that can be engaged as they hear instances of the term in conjunction with relevant exemplars, thereby acquiring the semantics of the term *on*.

This general hypothesis requires understanding three different aspects of acquisition: First, we need to know how infants and young children represent the notion of support before they learn language. A substantial literature has already shown that infants undergo considerable development in their non-linguistic understanding of support, starting with a rather coarse understanding of support, whereby an object is stable (i.e., will not fall) when half or more of its *bottom surface* is supported < ~ 12 months) but developing much richer and nuanced understanding by 13 months, when they come to understand the 'proportional-distribution rule', whereby an object is stable (will not fall) when half or more of the *entire object* is supported (see Baillargeon & DeJong, 2017). This might suggest that children's meaning of the term *on* changes as their understanding of physical support becomes more nuanced.

Second, we need to know how children represent the meaning of the spatial term *on*. Theories of spatial language meanings have often focused on the importance of so-called 'force-dynamic' relationships between objects, especially noting that terms like *in* and *on* require understanding how objects interact with each other to produce the relationship of 'containment' or 'support' (Vandeloise, 1991; Coventry et al., 2001). Understanding such force-dynamic relationships is required for mature use of terms *in* and *on* (at least in the case of relationships between concrete objects.) Several studies of young children's spontaneous production (Lakusta et al., 2020) and comprehension of *on* (Lakusta et al., 2021; Casasola & Park, 2013), suggest that children learning English map the basic locative construction '*is on*' to canonical examples of support, i.e., cases in which one solid object is supported from below by another solid object, as in "The apple is on the table" (Landau et al., 2017). The preference for encoding such instances of support-from-below with the basic locative construction of one's language (BE *on* in English; Levinson & Wilkins, 2006) has also been shown in children learning Greek, Hindi and Mandarin (Landau et al., 2017; Johannes

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

et al., 2016). In these studies, stimuli have portrayed 'core' exemplars of support with images of one object centrally positioned on top of another larger object, yielding a clear percept of support-from-below.

Although the bodies of data on prelinguistic understanding of support and later use of spatial expressions to describe 'core' exemplars make clear that the basic locative construction ('BE *on'*) is mapped onto canonical exemplars of 'support' (i.e. support-from-below by a solid body), we know little about how children's and adults' use of *on* maps onto their broader understanding of the intuitive physics of support under varying conditions of actual physical support. If the meaning of *on* entails understanding the force-dynamic conditions in which one object is supported by another, we might expect that children and adults would use 'is *on'* only when one object is fully supported by another (hence it will definitely not fall). Alternatively, uses of *on* could vary depending on the likelihood of one object being supported or not, with felicitous use of *on* ranging broadly over configurations that will definitely not fall and those that most likely or certainly will fall.

We test these possibilities by examining in detail the conditions under which young children and adults will respond to two different questions about the same set of spatial configurations which portray two objects, one of which (X) varies in degree of physical support from below by another (Y). The first question engages the use of language, i.e. the term *on*: "Is X 'on' Y?" The second is a question that engages one's intuitive physics: "Will X fall if (agent) lets go?" We ask whether the two questions will elicit the same pattern of responses. If the expression 'is on' is felicitous only when one object is truly supported by another, then we predict that the two questions should elicit heavily overlapping response patterns for the same configurations, i.e., if X is truly supported by Y, then the answers should be "Yes, it's *on*" and "No, it won't fall". Alternatively, if the two systems of knowledge (language, intuitive physics) map onto each other in only a partial and more complex way, then we might expect rather different response patterns to the two questions for the very same configurations.

At first glance, the question of whether one would agree that 'X is on Y' in just those circumstances where 'X will not fall' may seem to have an obvious answer– that is, yes. However, several bodies of literature on both sides of the equation suggest that the mapping between these two types of judgments (Is it on? Will it fall?) may not be straightforward and/or simple. For one thing, children's understanding of intuitive physics during early childhood is not sufficiently well-understood to provide a clear answer (Hood, 1995; Joh, Jaswal, & Keen, 2011) but we do know that even educated adults find certain judgements of force-dynamics challenging (McCloskey, Washburn, & Felch, 1983). Moreover, it is likely that use of '*is on'* by both children and adults varies substantially across configurations in which there are varying degrees of support for the target object. Here, we ask what the relationship is between use of *on* for the target object and likelihood of support as assessed by judgments of intuitive physics (here, whether the object will fall).

Method

Participants

Participants were 120 adults who were recruited via Prolific and 42 children aged 4 years (M=4 years 6 months (1679 days)) who were recruited through the LookIt platform. All participants were native English speakers and participants/parents of all participants signed informed consent before beginning the study. Both adults and children were tested online (Prolific for adults; Zoom for children).

Design and Materials

Design A mixed design was utilized with Condition (Language, Intuitive Physics) as a between-subjects factor and Degrees of Contact[1] (1-5) and Orientation (Regular/Mirror Image) of the target object as within-subjects factors. Participants were randomly assigned to the Language or Intuitive Physics condition, and each participant received one trial for each degree of contact (1-5) and orientation (Regular/Mirror Image), for a total of 10 trials per participant (see Figure 1). Trial types were presented in a different, random order for each participant.

Materials Stimuli consisted of videos depicting a puppet placing an L-shaped object so that some proportion of it was in contact with a table. The L-shaped object was presented in the Regular L-shape orientation as well as a Mirror Image L-shape orientation (Figure 1). For each orientation (Regular and Mirror image) the object was placed on the table at the five different degrees of contact. These were specifically chosen to include 1) a fully unsupported configuration (10% of entire object in contact with table) 2) a fully supported configuration (100% of the object in contact with the table) and 3) three configurations in which the object was supported if it was placed in the regular L shape orientation on the table, but not supported if it was placed in the mirror-image L shape orientation on the table (see Figure 1)[2].

[1] Degrees of contact corresponded to how much of the object's volume was supported by the table. These degrees of contact varied from 10% to 100% of volume supported.

[2] Following Bonawitz, et al. 2012, calculations were conducted to establish the exact (x,y) coordinate point of the L-shape's center of mass. After finding the object's center of mass, we were able to identify how much of the object's base needed to be supported from below by the table in order to not fall. Approximately 65% of the object's base must have been in contact with the table to afford full support.

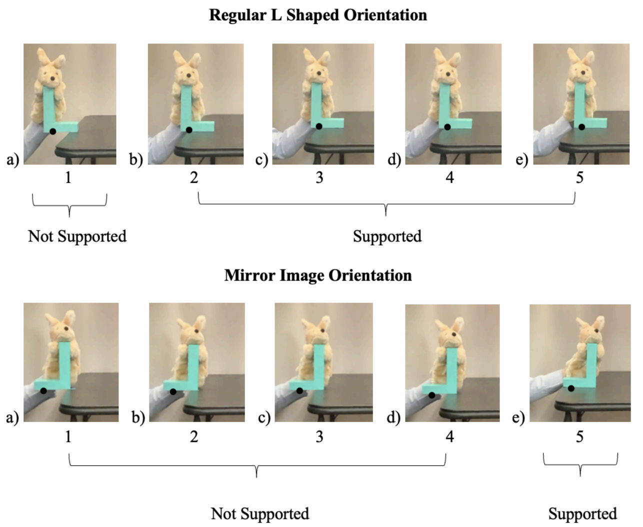

Figure 1: Placement of L-shaped object so at least 10% of it was in contact with a table; the L was placed in each of two orientations (Regular and Mirror Image). For each orientation, the L was placed such that the object achieved varying Degrees of Contact, ranging from full contact (5) to varying levels of contact depending on the orientation (1-4). Actual support was calculated by determining how much of the object's base needed to be supported (indicated by black dot, which was not visible to participants) using the center of mass rule (see Footnote 2). In the Regular orientation this resulted in no support for the object in panel a, and support for objects in panels b, c, d, and e. In the Mirror Image orientation, the placements resulted in no support for objects in panels a, b, c, and d, but full support for the object in panel e.

The videos were pre-recorded and standardized to ensure consistent visual and contextual presentation. The puppet's actions were designed to clearly show the object's placement and amount of contact provided by the table.

Procedure

To familiarize participants with the task and ensure comprehension, the experiment began with a training phase. Participants were introduced to the yes/no response format by first using simple questions unrelated to the task (e.g., "Is this a kangaroo?" pointing to an image that either was or was not a kangaroo). Participants were then introduced to two L-shaped objects (each named a "Riff") in two different orientations: a regular L-shaped orientation and its mirror image, i.e. "This is a Riff! Here it is on one side (depicted regular orientation) and here it is on its other side (depicted mirror image orientation). Today you are going to answer some questions about the Riff.". Participants were then introduced to a puppet character, Kara the Kangaroo, who demonstrated attempts to place the Riff 'on' a table and saw pictures of Kara the Kangaroo holding onto the L-shaped object in both orientations. Following these introductions, the task was to decide whether Kara successfully placed the Riff on the table.

Participants then started the second part of training. They observed two videos: one in which Kara the Kangaroo placed and held the Riff directly on top and in the middle of the table and another in which Kara the Kangaroo held the Riff completely off of the table. After each video, participants were asked the following questions depending on which condition they were assigned to. Those in the Language condition were asked "*Is the Riff on the table?*" and those in the Intuitive Physics condition were asked "*Will the Riff fall if Kara lets go?*". For both types of training trials, participants answered all pre-training questions correctly. This ensured that all participants fully understood the task requirements before advancing to the test phase.

During test trials, participants observed pre-recorded videos of the puppet placing the L-shaped object in the relevant positions (see Design and Materials and Figure 1). The puppet held the object in place and announced when they were finished by stating "I'm done!", which indicated the end of a trial.

In the Language condition, after the puppet placed the object, held it there and stated they were done, participants were asked, *"Is the Riff on the table?"*. In the Intuitive Physics condition, the procedure was identical except participants were asked, *"Will the Riff fall if Kara lets go?"*.

Children completed the study via Zoom, with an experimenter providing verbal instructions and guiding them through the trials. Adult participants completed the task independently through the platform Prolific, where they received written instructions and submitted their responses.

The total duration of the procedure was approximately 10 minutes, including training. Binary responses (*yes* or *no*) were recorded for each test trial.

Results

The raw data were binary coded with 1 indicating "yes" for the Language condition/ "no" for the Intuitive Physics condition. Recall, for the Language and Intuitive Physics conditions, respectively, participants were asked "Is the Riff on the table?" and "Will the Riff fall if Kara lets go?". Answering "yes" to the question "Is the Riff on the table?" in the Language condition and "no" to the question "Will the Riff fall if Kara lets go?" in the Intuitive Physics condition are *equivalent* in terms of predicting that the object will be supported; thus, both were coded as '1' for the analyses. These data were analyzed using a mixed-model logistic regression analysis. Age group (Adults, Children) and condition (Language, Intuitive Physics) were between subjects variables and degree of contact (1-2-3-4-5) and orientation (Regular/Mirror-Image) were within subject variables. These factors and all the interactions among them were entered as fixed effects, and participant was entered as a random effect.

The omnibus tests, which assess the overall effects of the predictors, yielded a significant effect of condition (Language, Intuitive Physics), X^2 (1) = 7.92, p = .005;

participants in the Language condition were more likely to answer, "yes" to the question, "Is the Riff on the table?" than participants in the Intuitive Physics condition were to answer "no" to the question, "Will the Riff fall…?". That is, as in Figure 2, responses in the Language condition, (green bars) are higher than responses in the Intuitive Physics condition (blue bars). This indicates that the Language and Intuitive Physics conditions elicit *different* patterns of responses (not showing complete overlap).

There was also a significant effect of degrees of contact, X^2 (4) = 19.30, $p < .001$. This indicates that participants were sensitive to increasing degrees of contact (1-2-3-4-5, see Figure 1), and were (correctly) more likely to judge that an object would not fall at greater degrees of contact.

The effects of group (Adults, Children) and orientation (Regular/ Mirror-Image) were not significant, *p*s >10.

In addition, there were two interactions among the factors. First, there was a significant three-way interaction among group (Adults, Children), condition (Language, Intuitive Physics), and degrees of contact (1-2-3-4-5), X^2 (4) = 11.32, p = .023. Second, we had anticipated that the two different orientations might elicit different patterns for adults and children (based on earlier work by Baillargeon & Hanko-Summers, 1990). We did find a marginally significant four-way interaction among the four factors group (Adults, Children) x condition (Language, Intuitive Physics) x degrees of contact (1-2-3-4-5) x orientation (Regular, Mirror Image), X^2 (4) = 9.01, p = .061.

As shown in Figures 2a and 2b, the pattern of responses differed across the two orientations of the L-shaped figure. Figure 2a shows that in the Language condition (green bars), a majority (above .50) of adults and children judged that the object 'is on' at all degrees of contact (1-5) but in the Intuitive Physics condition (blue bars), their responses were more graded - increasing as the degree of contact increased.

Figure 2b shows the parallel results for the Mirror Image orientation. In the Language condition (green bars), neither adults nor children were categorical (i..e, were not at ceiling across the degrees of contact as in the regular orientation). Rather, their responses were more graded - increasing as the degrees of contact increased. In the Intuitive Physics condition (blue bars), their responses were also more graded, and only at ceiling for objects shown at the highest degree of contact (i.e. at 5, when the object was supported).

All other two- and three-way interactions were not significant, *p*s > .10.

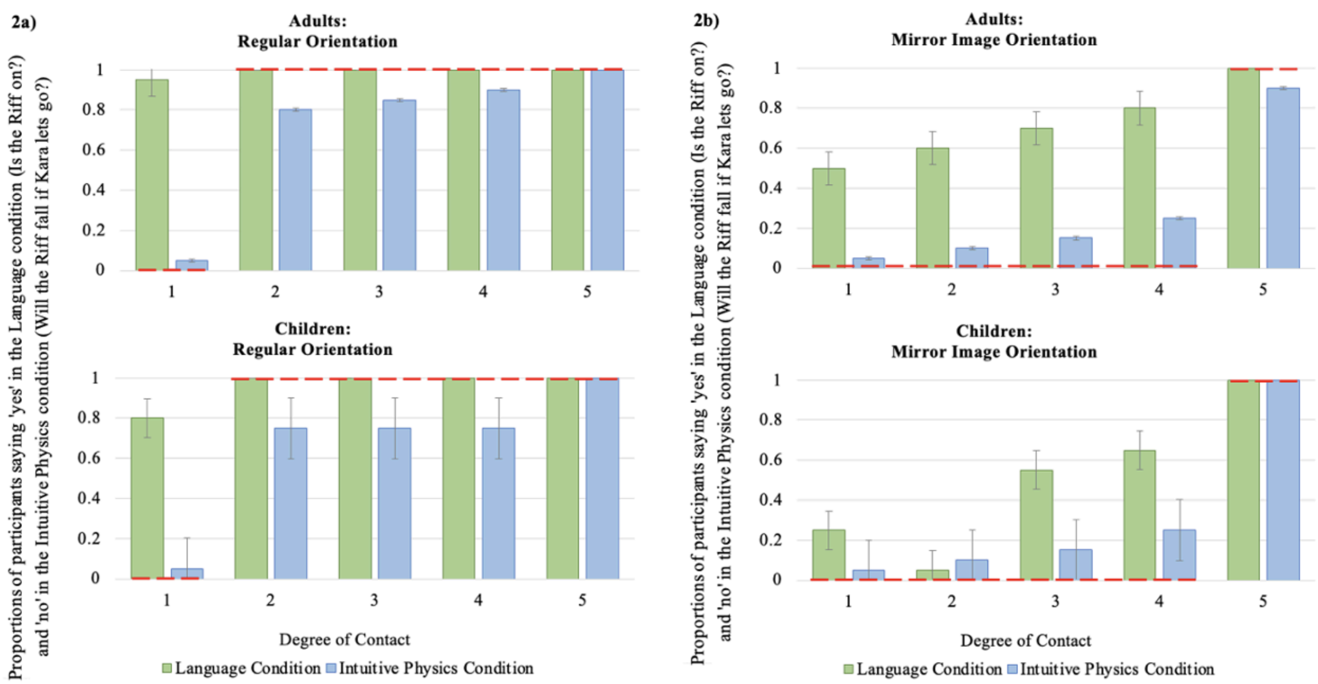

Figure 2: Bar graphs show the proportion of adult (top) and child (bottom) participants answering the Language ("Is the Riff on the table?", green bars) and Intuitive Physics ("Will the Riff fall…", blue bars) questions. Panel 2a (left) shows responses for the Regular L-shape orientation and Panel 2b (right) for the Mirror-Image orientation. The red dotted line indicates the likelihood of the object actually falling based on the law of center of mass. See text for summary of findings in these graphs.

Discussion

It is generally assumed that the acquisition and use of simple spatial terms (such as *in/on*) initially reflect rather simple mapping onto non-linguistic foundations. For the term *on*, the foundation is usually thought to be support– X is *on Y* if it is supported by Y (canonically, from a solid surface below). To test whether the meaning of *on* is directly related to the likelihood of support, we showed participants an L-shaped object being placed with varying degrees of contact relative to a table and we asked them one of two questions. The Language question was "Is the Riff on the table?"; the Intuitive Physics question was "Will the Riff fall if Kara lets go?". If the meaning of *on* reflects a simple mapping to intuitive physics representations of whether the object will fall, the response patterns should be consistent across the two questions. If so, this would yield evidence that uses of *on* by adults and children fundamentally reflect their computation of whether X is supported by Y. However, the findings showed that the two questions elicited quite different patterns of responses, suggesting that the two queries engage different domains of knowledge which are not completely coupled. That is, use of the basic spatial expression 'is on' does not always yield the same judgment as asking whether the supported object will fall. Thus, for this particular spatial term (*on*), language and intuitive physics reflect only a partial mapping; though deciding whether X is on Y may draw on intuitive physics in some/ many cases, there is no simple 1:1 relationship between the use of this spatial term and its putative non-linguistic foundation.

The details of our findings show that a full understanding of how simple spatial terms draw on non-linguistic foundations will require embracing some complexity. To exemplify, we structure our discussion in terms of two puzzles. One is why participants' judgments for the Language and Intuitive Physics questions were so different from each other for the same stimuli. As a first example, consider people's responses for the regularly oriented L-shaped objects. For both adults and children, responses in the Language condition were largely categorical; participants judged 'is on' for all degrees of contact, even with the least likelihood of true support (see Figure 2a, high green bars when degree of contact = 1). In the Intuitive Physics, by contrast, participants showed a graded - *relatively accurate* - pattern of responses; i.e., the blue bars in Figure 2a are near the red bar which represents true support of the object. Our finding that the spatial term *on* was used even in conditions where the object might (or would definitely) fall, shows clearly that participants were using *on* across a wide variety of scenarios portraying true support but also failures to support. As Talmy (1983) and others have noted, closed class terms as a whole do not engage metric information; details about precise location can always be added in a phrase (e.g., 'almost on', 'completely on', etc.). Intuitive physics, on the other hand, appears to naturally engage more precision in determining whether something will fall. These generalizations applied to the data from both adults and 4-year-old children. Note that it is an open question whether children are sensitive to more complex expressions that combine with *on* to produce less categorical responses. Indeed, if we had modulated our question of "Is the Riff on the table?" with additional phrases (e.g., Is it almost/completely/ partially on?) participants might have readily understood these modifications, and answered accordingly, yielding less categorical data than we got from our simpler question "Is the Riff on the table?".

The idea that the pragmatics of the simple "Is X on Y?" question may bias the participant to make categorical decisions is consistent with our understanding of how word meanings draw on non-linguistic foundations to determine rules of application. It is well-known that use of a lexical term does not map in detail to the structure of our non-linguistic systems. An obvious case is the color system, where the numbers of color terms in a language requires setting up lexical divisions that vary across languages (Berlin & Kay, 1969). Although we can discriminate many hues, not all hues are encoded lexically and part of the learning task is determining how the numbers of basic color terms in a language carve up the spectrum. Spatial terms are no different; there are far fewer basic spatial terms in a language than there are exact locations in space (Talmy, 1983). Judgments of whether or not an object will fall when it is placed in a given location relative to a (possibly) supporting object may be fine-grained enough to allow us to decide whether to place a bowl at the edge of a table; but this resolution is not needed to communicate the general location of the bowl.

The second remaining puzzle is why responses to the Language and Intuitive Physics questions differed from each other across the regular and mirror-image orientations. For the Language question, adults' and children's responses for the mirror image orientation suggests a more graded representation (relative to the regular orientation), with judgments that the L 'is on' steadily increasing from 1 to 5 degrees of contact (Figure 2b, green bars). The only exception was for children, who were conservative in saying "yes" to "Is the Riff on the table?" at both 1 and 2 degrees of contact. For the Intuitive Physics question, participants' responses for the mirror image orientation also showed a graded - *relatively accurate* - pattern of responding, as it did for the regular orientation (blue bars in Figure 2b are near the red bar which represents true support of the object). But note that, for the mirror-image orientation (Figure 2b), despite both language and intuitive physics suggesting graded representations, the response rates still differed. Language judgments about whether the object "is on" the table (green bars) were higher than intuitive physical judgments about whether something would not fall (blue bars). As in the regular orientation, participants were using *on* liberally.

The ability to make precise predictions about whether one object will fall without support from below becomes quite complex as we examine configurations that include objects with asymmetries. For asymmetrical objects, the entire object (not the bottom surface of the object) needs to be considered when determining whether an object will fall, and this understanding takes time to develop in infants (Baillargeon & DeJong, 2017). However, unlike infants, in the current study, intuitive physics reasoning was quite accurate for adults as well as children, even for the mirror image orientation. That is, in Figure 2b, intuitive physical reasoning among both children and adults tracks closely to the red line, which represents actual support of the object. What varied between the two orientations is how the word *on* was applied. Although *on* was used quite liberally by participants across both orientations, its use was more categorical for the regular orientation and more graded for the mirror image. To us, this suggests that understanding of the intuitive physics of complex objects may ultimately penetrate into the uses of words, in this case, the use of 'on' to indicate support.

To our knowledge, there are no studies that examine how "is on" maps to asymmetrical objects that are supported from below. One possibility is that as computation of support becomes more complex (e.g., for asymmetric objects whose parts fall outside of the supporting object's convex hull), people might use a simple rule similar to those used by infants as they develop more sophisticated ways of computing support. A different possibility is that the simple expression "is on" (i.e., the basic locative construction, Levinson & Wilkins, 2006) most readily maps to symmetrical objects that are supported from below (i.e., 'simple' cases of support). 'Harder' cases of support (e.g., those involving asymmetrical objects such as those used in the current study) may invite more complex linguistic descriptions, including lexical verbs and modulators such as those we mentioned earlier (e.g., almost on, just about on, etc.). Examining how people describe the different support relations we built into our stimuli might provide clues to how people represent 'support' across our stimuli. The use of simple locative expressions (e.g,. 'is in/on') for simple containment and support relations, and the extensive use of verbs and structural variation to capture more complex examples of containment and support relations has been shown in a variety of studies. For example, whereas support-from-below is most often described with the basic locative "is on", more complex configurations such as a coat on a hook or a stamp on an enveloped elicit a range of lexical verbs and prepositions ("hang from", "sticks to", etc.; see Landau, 2020 and Lakusta, Wefferling, Elgamal, & Landau, 2024). In order to explore whether our varying exemplars of support elicit different kinds of expressions, our next steps could include varying the shape and degree of asymmetry of the target object, setting up conditions that parallel the ones used by Baillargeon and colleagues, and determining whether the 'simple' cases elicit simpler expressions, while 'harder' cases elicit more complex linguistic structures.

In sum, the current study provides strong evidence that use of the spatial term *on* does not map simply to judgments of intuitive physics and that use of the spatial term *on* may engage, but is not determined by, whether or not a person believes the target might or might not fall. More generally, there are significant limitations in the extent to which an intuitive physics judgment will 'deliver' information that can be transparently used to decide whether X is *on* Y. What may have seemed to be a simple mapping is not so simple after all.

Acknowledgements

This research was supported by a National Science Foundation grant (1650861) awarded to Laura Lakusta and Barbara Landau.

References

Baillargeon, R., & Hanko-Summers, S. (1990). Is the top object adequately supported by the bottom object? Young infants' understanding of support relations. *Cognitive Development*, 5(1), 29–53. https://doi.org/10.1016/0885-2014(90)90011-H

Baillargeon, R., DeJong, G.F. Explanation-based learning in infancy. *Psychonomic Bulletin & Review*, 24, 1511–1526 (2017). https://doi.org/10.3758/s13423-017-1334-4.

Berlin, B., & Kay, P. (1969). *Basic Color Terms: Their Universality and Evolution*. Berkeley & Los Angeles: University of California Press.

Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*(4), 215–234. https://doi.org/10.1016/j.cogpsych.2011.12.002.

Casasola, M., & Park, Y. (2013). Developmental changes in infant spatial categorization: When more is best and when less is enough. *Child Development*, 84(3), 1004–1019. https://doi.org/10.1111/cdev.12010.

Coventry, K.R., Prat-Sala, M., & Richards, L. (2001). The interplay between geometry and function in the comprehension of over, under, above, and below. *Journal of Memory and Language*, *44* (3), 376, 398.

Hood, B. M. (1995). Gravity rules for 2-to 4-year olds?. *Cognitive Development*, 10(4), 577-598.

Joh, A. S., Jaswal, V. K., & Keen, R. (2011). Imagining a way out of the gravity bias: Preschoolers can visualize the solution to a spatial problem. *Child Development*, 82(3), 744-750.

Johannes, K., Wilson, C., & Landau, B. (2016). The importance of lexical verbs in the acquisition of spatial prepositions: The case of in and on. *Cognition*, 157, 174-189.

Lakusta, L., Brucato, M., & Landau, B. (2020). Evidence for a core representation for support in early language development. *Language Learning and Development*, 16(2), 180–195. https://doi.org/10.1080/15475441.2020.1721291.

Lakusta, L., Hussein, Y., Wodzinski, A., & Landau, B. (2021). The privileging of "support-from-below" in early spatial language acquisition. *Infant Behavior and Development*, 65. https://doi.org/10.1016/j.infbeh.2021.101616.

Lakusta, L., Wefferling, J., Elgamal, K., & Landau, B. (2024). "Dividing the labor": Lexical verbs and the linguistic encoding of physical support in 2- to 4.5-year-old children. *Journal of Experimental Child Psychology*, 238, 105803. https://doi.org/10.1016/j.jecp.2023.105803.

Landau, B. (2020). Learning simple spatial terms: Core and more. *Topics in Cognitive Science*, 91-114.

Landau, B., Johannes, K., Skordos, D., & Papafragou, A. (2017). Containment and Support: Core and Complexity in Spatial Language Learning. *Cognitive Science*, 41 Suppl 4, 748–779. https://doi.org/10.1111/cogs.12389.

Levinson, S. C., & Wilkins, D. P. (Eds.). (2006). *Grammars of space: Explorations in cognitive diversity.* Cambridge University Press. https://doi.org/10.1017/CBO9780511486753.

McCloskey, M., Washburn, A., & Felch, L. (1983). Intuitive physics: the straight-down belief. and its origin. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(4), 636.

Talmy, L. (1983). How Language Structures Space. In: Pick, H.L., Acredolo, L.P. (Eds.), *Spatial Orientation*. Springer, Boston, MA. https://doi.org/10.1007/978-1-4615-9325-6_11.

Vandeloise, C. (1991). *Spatial prepositions: A case study from French.* University of Chicago Press.