

Emergence of Higher-Order Transitivity across Development: The Importance of Local Task Difficulty.

Heidi Kloos

Department of Psychology, University of Cincinnati, Cincinnati OH, 45221

Abstract—This study investigates the effect of local task difficulty on children’s tendency to combine pieces of information into larger wholes. The particular hypothesis is that the emergence of higher-order Gestalts is guided neither by innate capabilities nor by laborious thought processes. Instead, it is – at least partly – tied to the difficulty of the local task, adaptively allowing the mind to reduce cognitive demand. The higher-order Gestalt used here was the transitive congruence among three feature relations. And the local task was to remember two of the feature relations from brief exposures, after having learned the third relation to criterion. The two to-be-learned relations violated transitivity with the third relation, such that a bias toward higher-order transitivity could be determined on the basis of children’s performance mistakes. Importantly, the two to-be-learned relations either matched in direction (posing a low cognitive demand), or they had opposite directions (posing higher cognitive demand). Results show that 5- to 9-year-olds were affected by higher-order transitivity only in the locally difficult task, not when the local task was easy.

Index Terms—Congruence, Feature Correlation, Gestalt, Task Difficult, Transitivity.

I. INTRODUCTION

A uniquely human ability is to interlink pieces of information into larger wholes. It lies at the heart of abstract thought, it allows us to form overarching theories that connect otherwise isolated elements, and it allows us to make inferences about unseen events on the basis of analogy. This ability to abstract higher-order patterns is present already in young children. For example, preschoolers will ignore individual features of an object in favor of higher-order category membership (e.g., Gelman, 1988). And they will ignore individual events that do not fit with their higher-order belief about cause-effect relations (e.g., Barrett et al., 1993; Karmiloff-Smith & Inhelder, 1975; Keil, 1989; Simons & Keil, 1995; Springer, 1995). What kind of process makes such pattern-seeking abstractions possible?

Two theories are commonly contrasted when it comes to understanding the development of higher-order thought. On the one hand, children might have an automatic and hard-wired tendency to combine individual features into higher-order patterns. For example, they might use innate ideas about

categories to classify novel items or infer hidden features (e.g., Gelman, 1990); they might apply abstract knowledge about causes and effects to interpret associations and contingencies (e.g., Opfer & Bulloch, 2007); and they might be guided by hard-wired assumptions about language (e.g., Golinkoff et al., 1994; Markman, 1989; Soja et al., 1991), social interactions (e.g., Tomasello, 1999) or physical truisms (e.g., Spelke et al., 1992) to organizing an otherwise messy input. Evidence for this view comes from numerous early-competence studies that look at young children’s dishabituation patterns, their decisions about category membership, their inferences and predictions, and their problem-solving performance. Common to these very different paradigms is the idea that children create higher-order patterns effortlessly, on the basis of powerful scaffolding mechanisms.

On the other hand, there is evidence suggesting that higher-order patterns require a laborious thought process, one that is difficult to attain by young children. Indeed, there are circumstances in which children seem surprisingly oblivious to abstract ideas about causes and effects (e.g., Kloos & Sloutsky, 2005); they ignore insights about higher-order category memberships (Sloutsky, Kloos, & Fisher, 2007); and they make blatant mistakes when it comes to thinking about higher-order physical principles (e.g., Hood et al. 2000). Indeed, children seem to have difficulty detecting salient contradictions in local pieces of information, whether the nature of contradictions is semantic (e.g., Osherson & Markman, 1975; Karmiloff-Smith, 1984), logical (e.g., Morris & Sloutsky, 2002; Ruffman, 1999) or inferential (e.g., Moshman & Franks, 1986). Taken together, these findings imply that the development of higher-order thought is slow, requiring mature processes of controlling one’s attention to organizing principles of Gestalt.

The discrepancy in points of view about what guides abstract thought can be illustrated with previous findings on transitive-inference tasks. The general procedure of such a task is to provide the learners with two feature relations and ask them to guess the third. For example, upon learning that the size of object A is larger than the size of object B, and that the size of object B is larger than the size of object C, the task is to guess the unseen size relation between objects A and C

(that the size of object A must be larger than the size of object C). There are circumstances in which children perform exceptionally well in such a task. For example, numerous categorization tasks show that – upon learning that an object with feature A also has feature B, and an object with feature B also has feature C – children easily assume that the object with feature A also has feature C.

On the other hand, traditional transitive-inference tasks are notoriously difficult for children (e.g., Riley & Trabasso, 1974). For example, upon learning that $A > B$, and $B > C$, children have difficulty making the transitive inference that $A > C$. Performance only improves after children participate in a lengthy learning phase where each individual relation is presented to them extensively, in both directions and across multiple sessions (e.g., Bryant & Trabasso, 1971). These findings imply that the emergence of transitive Gestalt is a laborious mental process, unattainable by young children.

The current study is designed to shed light on such discrepant findings. Why do children sometimes easily form higher-order patterns, while other times ignoring evident contradictions altogether? Answering this question has important implications, both for our understanding of the child's mind, and for pedagogical considerations geared towards learning. In particular, if higher-order thought is effortless and hard-wired, ideal instruction merely needs to support these natural-occurring tendencies. But if higher-order thought is laborious and hard to accomplish by young children, ideal instruction either provides extensive training in abstract reasoning or might spare young children altogether.

The general assumption put forward in this paper is that higher-order thought is a function of the actor-environment relation established in the immediate task context (cf., Stephen et al., in press). In other words, higher-order thought is neither automatic and effortless, nor controlled and effortful. Instead, it is an adaptive tendency to minimize the cognitive demands of a task. Thus, the reason why children sometimes focus on individual elements, and other times on higher-order patterns, should be found in the nature of the specific mind-task relations. The current paper tests this hypothesis by manipulating the difficulty of a local learning task.

A. Transitivity among Feature Correlations

Heavy objects tend to be big, a relation that can be represented as heavy = big (or $A = B$). Furthermore, heavy objects sink faster than light objects (with no changes in volume), a relation that can be represented as heavy = fast (or $A = C$). Finally, big objects sink more slowly than small objects (with no changes in mass), a relation that can be represented as big = slow (or $B = \neg C$). These three feature relations are the local elements of a higher-order pattern of transitive congruence (cf. Heider, 1958). For example, the relations mentioned above, heavy = big, heavy = fast, and big = slow violate logical transitivity: they are incongruent among each other. Figure 1A illustrates the incongruence of these relations graphically. In contrast, relations such as heavy = big, heavy = fast, and big = fast (as shown in Figure 1B) are an example of a congruent set of relations (relations that do

not violated transitivity).




Figure 1. Schematic representation of three feature relations among heaviness, size, and speed. They are either incongruent (A) or congruent (B) among each other.

This model domain of transitive feature correlations lends itself for investigating the emergence of higher-order Gestalt, for two reasons. First, previous findings have shown that young children can be sensitive to violation of transitive congruence among three feature relations (Kloos, 2007). For example, when presented with the set of relations shown in Figure 1A (heavy = big, heavy = fast, and big = slow), 4- to 7-year-olds were likely to misremember one of the relations as a means of obtaining a congruent set (e.g., they would mistakenly assume that big = fast, yielding the congruent set of relations shown in Figure 1B). Importantly, when the exact same sinking-speed relations were presented as part of a congruent set (i.e., heavy = small, heavy = fast, small = fast), children rarely made mistakes remembering them. The same two relations – when aligned with logical transitivity – yielded successful learning. These findings establish that young children can be sensitive to the higher-order Gestalt of transitive congruence among three relations.

Second, the higher-order configuration of three relations makes it possible to manipulate learning difficulty of local relations. For example, the two speed relations (mass-speed; volume-speed) could be the local relations that one is expected to learn from short demonstrations. And if these two relations have opposite directions (e.g., heavy = fast; big = slow), they would be more difficult to learn than when they match in direction (e.g., heavy = slow; big = slow). In the former case, when the two relations have opposite directions, one has to keep in mind that mass and speed correlate positively (heavy = fast), while volume and speed correlate negatively (big = slow). In other words, one has to attend to each feature relation separately – something that is likely to tax a child's cognitive system considerably. In contrast, when two relations match in direction, they can be collapsed and summarized under a simple rule of more = more (or more = less; cf. Smith & Sera, 1992). This learning manipulation makes it possible to manipulate the difficulty of the learning context independently of the higher-order structure of transitive congruence.

The current experiment builds on these considerations to test whether task difficulty predicts the emergence of higher-order Gestalt. Participants were asked to learn two sinking-speed relations that either matched in direction (matching-directions condition) or had opposite directions (opposite-directions condition). In both conditions, the two to-be-learned relations were embedded in an incongruent set of relations (see Figure 2). If children pay attention to higher-order

Gestalt, they should have difficulty learning the two feature relations. Conversely, if children focus on the local feature relations only, they should be unaffected by the lack of transitive congruence.




Figure 2. Schematic representation of three feature relations used in the opposite-directions condition (A) and the matching-directions condition (B). Note that the label for the condition pertains to the two speed relations.

There were three relations participants were exposed to: mass-volume, mass-speed, and volume-speed. However, the learning task emphasized only the two speed relations (shown with solid lines in Figure 2), by creating a challenge to learn them after only short demonstrations. Participants in the *opposite-directions* condition had to learn that heavy = fast and big = slow (Figure 2A); and participants in the *matching-directions* condition had to learn that heavy = slow and small = fast (Figure 2B). Note that the direction of the volume-speed relation is the same in both conditions (big = slow is equivalent to small = fast). Therefore, all children were asked to learn at least one negative feature correlation.

The relation shown with a dotted line pertains to the mass-volume correlation that renders the two to-be-learned relations incongruent. Participants were exposed to this relation in a preliminary phase. Rather than having to learn this relation from short demonstrations, they were given explicit instruction and feedback training about this relation. Despite yielding incongruence in both conditions, we predict that children should be affected by this initially learned mass-volume relation differently as a function of condition. In particular, we predict a stronger bias toward establishing higher-order transitivity in the opposite-directions condition than in the matching-directions condition.

II. METHOD

A. Participants

Children between 3 and 9 years of age were recruited from urban and suburban daycares and elementary schools serving predominantly middle-class families of mixed racial background. The final sample (61 boys and 59 girls) consisted of 21 3-year-olds ($M = 3.76$ years), 19 4-year-olds ($M = 4.47$ years), 19 5-year-olds ($M = 5.49$ years), 14 6-year-olds ($M = 6.48$ years), 16 7-year-olds ($M = 7.46$ years), 18 8-year-olds ($M = 8.43$ years), and 13 9-year-olds ($M = 9.18$ years). In addition, we tested 40 undergraduate students ($M = 20.4$ years; 23 men, 17 women), recruited from Introduction-to-Psychology courses, in return for class credit. Across children and adults, there were 80 participants per condition (60

children 20 adults). A total of 26 children were tested but excluded from the final sample because they did not meet the inclusion criterion (see Procedure). Most of these children were 3- to 4-years-old ($N = 19$), and, across age, most of these children were in the matching-directions condition ($N = 17$).

B. Materials


Stimuli consisted of real-life cylinders, pictures of the cylinders, and movies depicting the sinking behavior of the cylinders. In particular, there were eight real-life cylinders (5 cm in diameter; painted black) that differed in their height (6, 9, 13, or 19 cm) and weight (130, 200, 300, or 450g). Horizontal lines (each 4 cm long) were drawn on the lower part of a cylinder as a visual cue to its heaviness: two lines for 130 g, three lines for 200 g, four lines for 300 g, and six lines for 450 g. Cylinders were combined in two sets of four cylinders, one for each condition. To convey the heavy = big relation (opposite-directions condition), four cylinders were created for which mass and volume correlated positively. And to convey the heavy = small relation (matching-directions condition), four cylinders were created for which mass and volume correlated negatively. The top part of Table 1 provides a schematic of these two sets of four cylinders.

Various schematic pictures were used to test participants' beliefs about the learned relations (mass-volume, mass-speed, or volume-speed). They are described in detail in the Procedure, as each of the tests is introduced. In addition to real-life cylinders and schematic pictures of the cylinders, short movie clips of cylinders sinking in a water tank were used. For each clip, two cylinders appeared above a realistically drawn water tank, and upon releasing them into the tank, they sank in a semi-realistic way, such that one of the cylinders sank perceivably faster than the other. Cylinders in a clip either differed in mass (i.e., number of lines) or in volume (e.g., height). This made it possible to convey how mass and volume, each considered separately, affected sinking speed. Size and number of lines of cylinder cartoons matched those of the real objects. There were four movies for which the heavier cylinder sank fastest, four movies for which the lighter cylinder sank fastest, and four movies for which the smaller cylinder sank fastest.

Table 1. Schematic depiction of the procedure (in three phases) and corresponding materials, separated by condition.

Task	Conditions	
	Opposite Direction	Matching Direction
1. Mass-Volume Correlation		
Training	Participants experience the mass/volume correlation	
Testing		
Hidden-Mass	"Which one is heavier?"	
Hidden-Volume	"Which one is taller?"	

2. Sinking-Speed Correlations	
Training	
Mass-Sinking	Participants see short movie demonstrations of sinking objects. Long arrow = fast.
Volume-Sinking	
Testing	
Mass Trials	“Which of the two cylinders to the left and right of S sinks faster than (the standard) S?”
Volume Trials	
3. Mass-Volume Correlation	
Testing	same as above



Note. Conditions differ in the mass-volume correlation and the mass-speed relation. All testing events and the volume-speed relations are the same across conditions.

C. Procedure

Participants were tested one-on-one with an experimenter, either in a quiet area at their school or in the lab. The cover story involved Scuba Sam, a character who builds submarines. The task was to find out about what makes submarines sink fastest in water, in order to win a submarine race against Scuba Sam. The experiment was administered with Superlab Version 2.0 on a Dell Laptop computer. First, participants were introduced to a single cylinder with four lines displayed on the computer. They were told: “Look this is a submarine of the submarine man. The lines show the number of weights inside the submarine.”

Table 1 shows the steps of the subsequent procedure in schematic form, separated by condition. It consisted of training and testing of the mass-volume correlation (Phase 1), training and testing of the two sinking-speed relations (mass-speed; volume-speed; Phase 2), and testing of the mass-volume correlation (Phase 3). Each of these phases is described in detail below.

Mass-Volume Correlation (Phase 1). Participants were first shown the set of four cylinders (depending on condition, either the positive-correlation set or the negative-correlation set). As a means of familiarizing participants with the cylinders, they were asked to order them according to mass (“Which submarine is the heaviest? Which one comes next?”), and then according to volume (“Which submarine is the tallest? Which one comes next?”). Help was provided as needed, in which case the ordering was repeated. After each ordering, the experimenter explained again the meaning of the lines (“See, the heaviest has the most lines”). The correlation between mass and volume was then explicitly pointed out (e.g., “See how the tallest one is the heaviest?”).

After cylinders were correctly ordered (according to mass and according to volume), they were removed from participants’ view. Mass-volume testing followed on the computer. During a hidden-mass trial, two cylinder pictures were displayed in such a way that the number of lines was

concealed. Participants were asked to determine the heavier of the two submarines, taking into account how tall the submarines are. And during a hidden-volume trial, two cylinder pictures were displayed in such a way that the height of the submarines was concealed. Participants were asked to determine the taller of the two submarines, taking into account the heaviness of the submarine (shown as number of lines). Participants completed four trials of each type, presented in blocked order. Feedback was provided after each trial, with each incorrect trial being replaced with a new trial, for a maximum of six trials per trial type. Participants had to perform correctly in at least four trials of each type to be included in the final sample.

Sinking-Speed Relations (Phase 2). Movie demonstrations of sinking objects followed next. For each movie, attention was first drawn to how the two submarines differed from each other. In particular, on mass-sinking trials, participants were told: “These two submarines have the same size, but one is heavier than the other. Which one is heavier?” And on volume-sinking trials, participants were told: “These two submarines weigh the same, but one is taller than the other. Which one?” Help was provided if needed. Participants were then instructed to watch closely to determine the winner of the race. After each race, the participant pointed to the submarine that sank faster. Feedback was provided (e.g., “Yes, the heavier one sank the fastest”). The movie was repeated when the winner was not detected correctly. Four unique mass-sinking movies and four unique volume-sinking movies were presented in random order.

Sinking-speed testing followed immediately to measure participants’ learning from the movie demonstrations. For each trial, participants were shown three cylinder pictures next to each other, with the middle picture serving as a standard. For mass trials, the pictures to the left and right of the standard differed only in the number of lines (they had the same volume as the standard). And for volume trials, the pictures to the left and right of the standard differed only in height (they had the same mass as the standard). The instruction was: “Now you get to race against Scuba Sam. This pink submarine in the middle of the screen is Scuba Sam’s submarine, and these two on the ends are yours. Which one of your submarines do you think would sink faster than Scuba Sam’s submarine?” There were twelve trials given in random order, six of which were mass trials (only mass varied), and six of which were volume trials (only volume was varied). No feedback was provided.

Mass-Volume Correlation Retested (Phase 3). The final phase consisted of another set of mass-volume testing trials to determine the extent to which participants could remember the mass-volume correlation present in the four cylinders shown at the beginning of the experiment. The testing trials were identical to mass-volume testing described above, with the following exceptions: four hidden-mass trials and four hidden-volume trials were presented in random order (rather than blocked); no feedback was provided; and no trials were repeated.

III. RESULTS

If considering isolated relations only, recall that the conditions differed in two relations: the direction of the mass-volume correlation, and the direction of the mass-speed relation. In preliminary steps, we looked at (1) how well participants could learn/remember the positive vs. negative mass-volume correlation; and (2) how well they could remember the positive vs. negative mass-speed correlation.




Figure 3. Mean proportion of correct mass-volume performance, separated by direction and phase. Standard errors are shown as error bars, and chance performance is 0.5.

Figure 3 shows the proportion of correct hidden-mass and hidden-volume trials (out of [the first] 8), as a function of age group, direction of correlation (positive vs. negative), and testing phase (Phase 1 vs. 3). It shows above chance performance in Phase 1 across ages, whether the mass-volume correlation was positive or negative. And in Phase 3, while performance was lower for the negative than the positive mass-volume correlation, average performance was above chance for all but the two youngest age groups. That is to say, efforts to establish incongruence among the three relations were successful for a large majority of participants older than 4 years of age.

In terms of the second preliminary question, Figure 4 shows the proportion of correct mass trials as a function of age group and condition (Phase 2). For comparison purposes, performance on volume trials is also provided (collapsed across conditions because it was negative for both conditions). As can be seen in the figure, there is no apparent developmental trajectory: means of all age groups but the 9-year-olds were at or above chance. And there was no main effect of trial type: whether learning pertained to a positive mass-speed relation, a negative mass-speed relation, or the negative volume-speed relation, mean performance was around 75% across age groups.




Figure 4. Mean proportion of correct performance on sinking-speed trials (Phase 2), separated by trial type (mass vs. volume) and direction (positive vs. negative).

For the main question – namely whether learning of the two sinking-speed relations was affected by condition – a congruence score was determined for each participant. It ranged from 0 (unaffected by congruence) to 6 (strongly affected by congruence). Performing correctly on all mass trials and all volume trials yielded a congruence score of zero (because the presented speed relations were incongruent with the mass-volume correlation). Therefore, the person who could learn both speed relations perfectly was unaffected by congruence. Conversely, performing correctly on all trials of one type, and incorrectly on all trials of the other type, yielded the maximum congruence score of 6 (because systematically misremembering one of the relations results in a congruent set of beliefs). In other words, the person who could learn one of the relations and systematically misremembers the other was highly affected by congruence.

Intermediate congruence scores were distributed accordingly (e.g., performing correctly on all mass trials and all but one volume trial received a congruence score of 1; and performing correctly on all mass trials and only one volume trial received a congruence score of 5). Participants who performed correctly on only half of the trials of one type did not receive a congruence score (which was the case for 38% of 3-year-olds, 16% of 4-year-olds, 12% of 7-year-olds, and 5% of each the 5-year-olds, the 9-year-olds, and the adults).

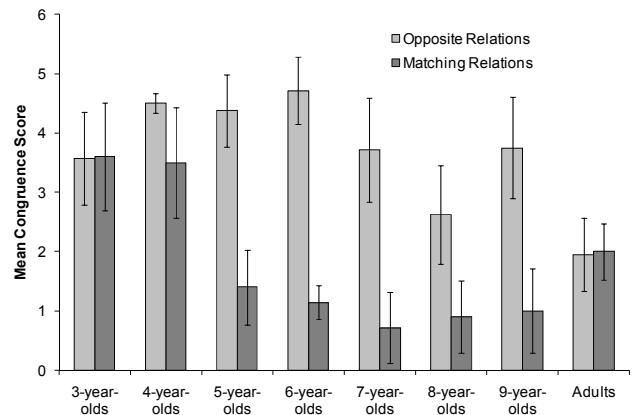


Figure 5. Mean congruence score as a function of condition and age group. Error bars represent standard errors.

Figure 5 shows the distribution of mean congruence scores as a function of condition and age group. As predicted, a

between-subjects ANOVA revealed a main effect of condition, $F(1, 150) = 5.7$, $p < .01$, with higher mean congruence scores in the opposite-relations condition ($M = 3.36$) than in the matching-relations condition ($M = 1.71$)¹. A significant condition by age interaction, $F(7, 150) = 3.1$, $p < .05$, stemmed from the fact that the effect of condition was strong for 5- to 9-year-olds, $t_s > 2.7$, $p_s < .05$, but lacking for 3-year-olds, 4-year-olds, and adults, $t_s < 1.1$.

IV. CONCLUSION

Children between 3 and 9 years of age had to learn two feature relations that either matched in direction (both were negative), or had opposite directions (one was positive, and the other was negative). This manipulation made it possible to vary the difficulty of learning local information (with opposite relations being more difficult to learn than two matching relations). Importantly, the two to-be-learned relations lacked transitive congruence with a third relation, which children learned to criterion in a preliminary phase. Adults were included to gauge of the developmental endpoint; and indeed they could learn the relations in both conditions – unaffected by the difficulty manipulation, and therefore unaffected by the overall incongruence. Conversely, 5- to 9-year-olds were highly affected by the difficulty manipulation. While they could learn the two relations when the relations matched in direction, they made systematic mistakes when the relations had opposite directions. That is to say, children in the latter condition sought to establish higher-order transitive congruence.

One could argue that children's performance in the opposite-relations condition merely reflects task difficulty, not an adjustment to higher-order congruence. To counter this argument, recall that 4- to 7-year-olds can easily learn two correlations that have opposite directions – as long as the two relations are couched within a congruent set (Kloos, 2007). The apparent mistakes in learning were instead a reflection of incongruence: children were trying to establish higher-order congruence when learning local elements were too taxing for the. This finding lends support the idea that the emergence of higher-order Gestalt –transitive congruence in this case – is neither entirely automatic nor entirely controlled. Instead it is related to task difficulty in an adaptive way (see also Adams, 1978). Parallel findings were reported with a gear task (Stephen et al., in press), and in children's perception of higher-order patterns (Kimchi, 1990). It remains to be determined how this adaptive feature of the mind can be harnessed for learning.

ACKNOWLEDGMENT

We thank Chris Erb, Kate Greene, Kellana Hindert, Chris Erb, and Sarah Schwind for their help with materials and data collection. Funding for this research was provided by NSF DRL 723638 and NICHD HD055324.

REFERENCES

- [1] Adams, M. J. (1978). Logical competence and transitive inference in young children. *Journal of Experimental Child Psychology*, 25, 477-489.
- [2] Barrett, S. E., Abdi, H., Murphy, G. L. & Gallagher, J. M. (1993). Theory based correlations and their role in children's concepts. *Child Development*, 64, 1595-1616.
- [3] Bryant, P. E., & Trabasso, T. (1971). Transitive inferences and memory in young children. *Nature*, 232, 456-458.
- [4] Gelman, S. A. (1988). The development of induction within natural kind and artifact categories. *Cognitive Psychology*, 20, 65–95.
- [5] Gelman, R. (1990). Structural constraints on cognitive development: Introduction to a special issue of *Cognitive Science*. *Cognitive Science*, 14, 3–10.
- [6] Gelman, S. A., & Markman, E. (1986). Categories and induction in young children. *Cognition*, 23, 183–209.
- [7] Golinkoff, R. M., Mervis, C. B., & Hirsh-Pasek, K. (1994). Early object labels: The case for a developmental lexical principles framework. *Journal of Child Language*, 21, 125–155.
- [8] Heider, F. (1958) *The psychology of interpersonal relations*. New York: John Wiley & Sons, Inc.
- [9] Hood, B. M., Carey, S., & Prasada, S. (2000). Predicting the outcomes of physical events: two-year-olds fail to reveal knowledge of solidity and support. *Child Development*, 71(6), 1540-1554.
- [10] Karmiloff-Smith, A. & Inhelder, B. (1975). If you want to go ahead, get a theory. *Cognition*, 3, 95-212.
- [11] Karmiloff-Smith, A. (1984). Children's problem solving. In M. E. Lamb, A. L. Brown, & B. Rogoff (Eds.), *Advances in Developmental Psychology*, 3, 39-90.
- [12] Kimchi, R. (1990). Children's perceptual organization of hierarchical visual patterns. *European Journal of Cognitive Psychology*, 2, 133–149.
- [13] Kloos, H. (2007). Interlinking physical beliefs: Children's bias towards logical congruence. *Cognition*, 103, 227-252.
- [14] Mandler, J. B., Bauer, P. J., & McDonough, L. (1991). Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology*, 23, 263–298.
- [15] Morris, B. J., & Sloutsky, V. M. (2002) Children's solutions of logical versus empirical problems: What's missing and what develops? *Cognitive Development*, 16, 907-928.
- [16] Moshman, D., & Franks, B. (1986). Development of the concept of inferential validity. *Child Development*, 57, 153–165.
- [17] Opfer, J. E., & Bulloch, M. J. (2007). Causal relations drive young children's induction, naming, and categorization. *Cognition*, 105, 206–217.
- [18] Osherson, D., & Markman, E. (1975). Language and the ability to evaluate contradictions and tautologies. *Cognition*, 3, 213–226.
- [19] Riley, C. A., & Trabasso, T. (1974). Comparatives, logical structures, and encoding in a transitive inference task. *Journal of Experimental Child Psychology*, 17, 187-203.
- [20] Ruffman, T. (1999). Children's understanding of logical inconsistency. *Child Development*, 70, 872–886.
- [21] Smith, L. B. and Sera, M. D. (1992). A developmental analysis of the polar structure of dimensions. *Cognitive Psychology*, 24, 99-142.
- [22] Simons, D., & Keil, F. C. (1995). An abstract to concrete shift in cognitive development: The insides story. *Cognition*, 56, 129-163.
- [23] Sloutsky, V. M., Kloos, H., & Fisher, A. V. (2007). When looks are everything: Appearance similarity versus kind information in early induction. *Psychological Science*, 18, 179–185.
- [24] Soja, N., Carey, S., & Spelke, E. (1991). Ontological categories guide young children's inductions of word meanings: Object terms and substance terms. *Cognition*, 38, 179–211.
- [25] Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, 99, 605–632.
- [26] Springer, K. (1995). Acquiring a naive theory of kinship through inference. *Child Development*, 66, 547-558.
- [27] Stephen, D. G., Dixon, J. A., & Isenhower, R. W. (in press). Dynamics of representational change: Entropy, action, and cognition. *Journal of Experimental Psychology: Human Perception and Performance*.
- [28] Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press.

¹ This difference held up even when we excluded participants who performed below 60% correct during the mass-volume testing in Phase 3.