Ambiguity and Representational Stability:
What is the role of embodied experiences?

Ramón D. Castillo (racastillo@utalca.cl)
Centro de Investigación en Ciencias Cognitivas, Universidad de Talca
Avenida Lircay s/n, Talca, CP 3460000, VII Región, Chile

Talia L. Waltzer (twaltzer@ucsc.edu)
Department of Psychology, University of California, Santa Cruz
1156 High Street, Santa Cruz, CA 95064 USA

Heidi Kloos (heidi.kloos@uc.edu)
Center for Cognition, Action, and Perception, Department of Psychology, University of Cincinnati
5130 Edwards One, Cincinnati, OH 45221-0376 USA

Abstract

Embodied cognition is sometimes presented as an alternative to computational approaches, the argument being that cognition is strongly influenced by an agent's body movement. However, the exact nature of this influence is still uncertain. In the current paper, we add to the conversation by analyzing adults' predictions in a high-ambiguity task: Adults had to decide which of two objects would sink faster (or slower) in water. Ambiguity was achieved by pitting object volume and object mass against buoyancy: The winning object of a pair was sometimes the bigger and heavier one, and sometimes it was the smaller and lighter one. The crucial manipulation was whether the stimuli were real-life objects or 2D pictures. All participants were presented with pictures of the objects during a training phase (when they received feedback on their predictions). Real-life objects were either present during the phase prior to the training (jars-first condition), or during the phase after the training (jars-last condition). Findings showed a clear influence of hands-on experiences: When allowed to hold the objects, adults were more likely to demonstrate a simplistic focus on object heaviness. These results call for a more nuanced understanding of the effect of embodied experiences on the stability of representations. While embodiment sometimes can help distinguish relevant from irrelevant information, we show that it can also destabilize representations acquired through visual information.

Keywords: action; knowledge representation; predictions; ambiguity; misconceptions; hands-on explorations

Introduction

What is the source of our thoughts, beliefs, attitudes, and the like? Traditionally, this question has been addressed with models of symbolic activity of the mind: Thoughts might be formed on the basis of combining symbols, which themselves are computed on the basis of simpler symbols, derived from sub-symbolic codes of sensation and perception. Approaches of embodied cognition stand in sharp contrast to the traditional view of computational models. They claim that mental activity, seemingly a bodiless manipulation of symbols, is instead strongly influenced by the very physical non-symbolic movement of our bodies (e.g., Chemero, 2011; Gibbs, 2005; Wilson & Clark, 2009). Rather than suggesting purely symbolic activities of bodiless minds, proponents of embodied cognition make a convincing case that higher-level cognition is constrained by our bodily experience of being in the world (Goldin-Meadow, Cook, & Mitchell, 2009).

In the current paper, we seek to explore the influence of embodied experiences in more detail. Our guiding theoretical framework does not subscribe to a specific representational format or cognitive architecture. Instead, we postulate that the mind makes use of whatever constraints are available in order to perform systematically (e.g., Kloos, Fisher, & Van Orden, 2010). These constraints could come from symbolic content, from bodily experiences, or from constraints outside the mental or bodily activity. The central question, then, pertains to how these different constraints interact. For example, to what extent does embodied experience override, support, or interfere with visual perception?

To explore this question, we analyzed the responses of adults in a high-ambiguity prediction task: Adults had to decide which of two objects would sink faster (or slower) in water. Ambiguity was achieved by pitting object volume and object mass against buoyancy: The winning object of a pair was sometimes the bigger and heavier one, and sometimes it was the smaller and lighter one. Thus, the task could not be solved with a simplistic rule that focuses on one dimension only (i.e., just weight or just size). To be successful, one must integrate both mass and volume by paying attention to the distribution of mass. While this integration can be accomplished, even by children (Kohn, 1093), it is not likely to be an adult's first guess (Castillo & Kloos, 2013). In fact, the initial tendency might be to focus on mass exclusively to make a decision (Castillo, Kloos, Richardson & Waltzer, 2015).

Note that high-ambiguity tasks, while not necessarily common in adults' everyday experiences, have been used extensively to better understand the mind's inner workings. The idea is that a high-ambiguity context reveals internal biases, natural preferences of the mind, so to speak. They are particularly useful to explore the role of embodied experiences: If embodied experience matters, then it should help disambiguate the constraints of the task. An added twist

here is that the specific task we chose – predicting the sinking behavior of objects – yields common misconceptions, namely that a reliance on mass alone could lead to successful performance. Would haptic explorations allow adults to overcome these misconceptions faster? Or would it in fact be more difficult for adults to benefit from such experiences?

Our overall method was as follows: Adults were presented with pairs of transparent objects that differed in size and contained a certain number of weights, clearly visible to participants. There were three phases: a *pre-test*, a *training*, and a *post-test*. Each phase had the same prediction trials, the difference being only in whether participants received feedback (training) or not (pre-test, post-test). The crucial manipulation was, before each prediction, whether adults were presented with real-life objects, or whether they merely saw the objects via 2D pictures. Specifically, one group of participants could explore real-life objects during the pre-test (jars-first condition), and one group of participants could explore real-life objects during the post-test (jars-last condition). During all other phases, stimuli were the 2D pictures of the objects. To what extent does the embodied experience affect performance?

Method

Participants

Participants were 112 adults between 18 and 27 years of age, recruited from a Midwestern university. They were each assigned to one of two conditions: the jars-first condition (17 men, 38 women; $M = 19.03$ years, $SD = 1.69$), or the jars-last condition (14 men, 43 women; $M = 19.02$ years, $SD = 1.67$). They received partial course credit for participation, following an IRB-approved procedure.

Material and Apparatus

Real-life sinking objects were used in this experiment, dropped into a water tank to create feedback for adults' predictions. The objects were transparent glass jars that differed in their sizes. Round aluminum discs (43g) could be placed inside the jars to manipulate mass. The water tank was 1m tall and had a vertical dividing wall to make it possible for each jar to sink without being affected by the other's turbulences.

The jars were combined into pairs of objects. Figure 1 shows an example for various different trials, which differ in how mass and volume correlated with rate of sinking. The faster sinking object within a pair is marked with a star. In two of the trial types, only one of the features was varied (either mass or volume), and in three of the trial types, both mass and volume were varied. Specifically, in the *small-wins* pair (Fig. 1A), mass was held constant and the size of the jar was varied in such a way that the smaller jar sank faster. In the *heavy-wins* pair (Fig. 1B), volume was held constant and mass was varied in such a way that the heavier jar sank faster. In the *big/heavy-wins* pair (Fig. 1C), the faster sinking object was bigger and heavier than the slower object. In the *small/heavy-wins* pair (Fig. 1D), the faster sinking object was

smaller and heavier than the slower object. And finally, in the *small/light-wins* pair (Fig. 1E), the faster sinking object was smaller and lighter than the slower object. There were nine pairs of each trial type, resulting in a total of 45 unique pairs.


Figure 1: Example trials. Star marks jar that sinks faster. A: Small-wins pair. B: Heavy-wins pair. C: Big/heavy-wins pair. D: Small/heavy-wins pair. E: Small/light-wins pair.

For simplicity, we will only report performance on the big/heavy-wins pairs (Fig. 1C) and the small/light-wins pairs (Fig. 1E). These two types of pairs create the high-ambiguity task context needed for the current purposes. This is because, while mass and volume correlate positively (the bigger of the two objects was also the heavier one), it was sometimes the heavier and sometimes the lighter object that sank fastest. Thus, to perform correctly, it would not be sufficient to pay attention to either mass or volume alone. All other trials had low ambiguity and will be considered fillers (indeed, adults performed largely at ceiling during those trials).

Pairs were presented either as actual jars or as pictures on a screen. Figure 2 shows the picture versions of the stimuli. Each trial included a close-up picture of a pair with discs outside the jar (Fig. 2A) as well as a close-up with discs inside the jar (Fig. 2B). Feedback was always provided as a picture of the jars being dropped in the tank of water (Fig. 2C). A numeric keypad was used to record participants' predictions.


Figure 2: Example pictures. A: Empty jars with weights on either side. B: Jars filled with weights. C: Jars sinking in the water tank.

Procedure

Participants were tested individually in the lab, using DirectRT Precision Timing Software (2012 Version) to administer the experiment on a desktop computer. The experiment consisted of a total of 360 prediction trials, divided into three phases. The first phase was the pre-test (90 trials): Participants made predictions across various jar combinations, without receiving any feedback. The second phase served as training (180 trials): Adults' predictions were followed by corrective feedback. Finally, the last phase was

the post-test (90 trials), featuring prediction trials that were identical to the pre-test (no feedback provided).

Our main manipulation was the timing of the embodied experience. Participants held the real objects either during the pre-test (jars-first condition), or during the post-test (jars-last condition). The training was always carried out with pictures. For generalizability purposes, we also manipulated the type of predictions adults had to make: Participants were either asked to predict which of two jars would sink faster, or which of two jars would sink slower.

During familiarization, participants were shown empty jars of different sizes, as well as several aluminum discs. They were told that all the discs have the same weight. The experimenter then filled the large and small jars with aluminum discs and asked the participant to predict which of them would sink faster in water (or slower). Participants were encouraged to lift the jars before making their predictions. Then they were provided with feedback pictures showing the outcome of the jars after being dropped in water. Finally, participants were shown the keypad and how it works. Prior to the experiment proper, they were informed that the pictures were taken from the real objects.

For predictions with real-life objects (pre-test or post-test, depending on condition), participants sat in front of an opaque box (60 x 25 x 40 cm) that served as a table to hold the objects. It also served as a barrier behind which the researcher kept the 12 jars (see Fig. 3 for a schematic over-head view of this arrangement). Based on a random order determined for each participant prior to the start of the experiment, the jar pairs were presented one at a time. For each pair, participants were asked to make a prediction about which one of the two jars would sink faster (or slower) in water. Participants were encouraged to respond by saying "left" or "right", corresponding to whether the winning (or losing) jar was in their left hand or right hand.


Figure 3: Diagram of the set-up for prediction trials with real jars. *R*: Researcher. *C*: Video camera. *P*: Participant.

For predictions with pictures (pre-test or post-test, depending on condition), participants were first shown an image of two empty jars next to each other, with a stack of discs by each jar. This allowed participants a clear view of the number of discs for each object. After 1.5 seconds, the image was replaced with a picture of the same two jars, but now filled with the discs and closed with a lid. Participants were asked to decide which of the two jars would sink faster (or slower). Figure 4A shows such a trial in schematic form.

There was no time restriction for making a prediction. The trial ended when the participant pressed the keypad to provide a prediction. A fifth of the all trials were big/heavy-wins trials, and a fifth of the trials were small/light-wins trials, interspersed with filler trials.

Training was identical to pre- and post-test predictions, the only difference being that a feedback picture was shown for 1.5 seconds, right after the participant made a prediction. On the very first feedback trial, the image was explained. The faster sinking object was pointed out on the computer screen, and participants were provided with explicit feedback (e.g., "Yes, you were right"; "No, look, it was the other one that sank faster"). Training took place between pre- and post-test. Of all the training trials, a fifth were big/heavy-wins trials, and a fifth were small/light-wins trials, interspersed with filler trials.
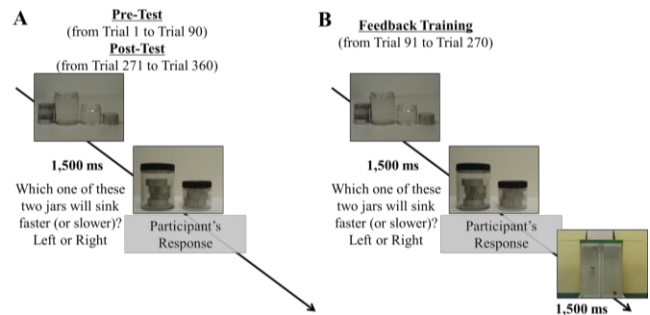

Figure 4: Schematic representation of the prediction trials.
A: Example picture trial during pre- or post-test.
B: Example picture trial during feedback training.

Results and Discussion

Our dependent variable was the proportion of correct predictions on big/heavy-wins and small/light-wins trials. Figure 5 presents the accuracy data for these two trial types, separated by phase (pre-test, training and post-test), and by the embodiment manipulation (jars-first vs. jars-last).


Figure 5: Proportion of correct answers for trial type and phase, separated by condition. Error bars represent standard errors of the mean. Circles highlight when real-life jars were used.

A 2 x 2 x 3 mixed-design ANOVA was carried out, with trial type (big/heavy-wins; small/light-wins) and phase (pre-test; training; post-test) as within-group factors, and with

condition (jar-first; jar-last) as the between-group factor. We found a significant 3-way interaction, $F(2, 220) = 22.47$, $p < .001$, $\eta^2 = .17$, prompting us to look at the results separately by phase (see Table 1 for a summary of the results).

During the pre-test, participants' performance was markedly different for the two types of trials: While performance was at ceiling (or close to) on big/heavy-wins trials ($M_{Pic} = .94$; $M_{Jar} = .94$), participants made systematic mistakes on the small/light-wins trials ($M_{Pic} = .37$; $M_{Jar} = .24$). It appears that participants resolved the ambiguity of the prediction task by focusing on weight exclusively. A 2 x 2 mixed-design ANOVA, with trial type and condition as factors, revealed a main effect of trial type, $F(1,110) > 100$; $p < .001$, a main effect of condition, $F(1,110) = 8.10$; $p < .01$, as well as a significant interaction, $F(1,110) = 4.74$; $p = .032$. The interaction is driven by the fact that participants' mistakes on small/heavy-wins trials were even more pronounced when they handled jars (jars-first condition) than when they viewed pictures (jars-last condition), $F(1,110) = 6.32$; $p = .013$.

During the training, the difference between trial types disappeared, whether participants were in the jars-first condition ($M_{Big/heavy} = .80$, $M_{Small/light} = 0.77$) or in the jars-last condition ($M_{Big/heavy} = .80$; $M_{Small/light} = 0.82$). A 2 x 2 mixed-design ANOVA (trial type by condition) yielded no main effects and no interaction, $Fs(1,110) < 2.37$; $ps > .12$. Performance was clearly above chance, $t(111) > 5$; $p < .01$, implying that adults benefited from the training and quickly learned that a focus on mass or volume alone yields mistakes.

To compare performance during training and pre-test, we carried out two 2 x 2 repeated-measure ANOVAs (trial type by phase), one for the jar-first condition, and one for the jar-last condition. For both conditions, the analysis yielded highly significant main effects and interactions, $Fs > 80$, $ps < .001$. While performance on big/heavy-wins pairs decreased slightly from pre-test to training in both conditions, $ps < .001$ it starkly improved for small/light-wins pairs, $ps < .001$. The results show that training had very similar effects on performance, whether participants had a chance to haptically explore the objects prior to training or not. This confirms that the switch between 3D objects to 2D pictures of the objects did not have a discernable effect on performance.

Finally, during the post-test, the difference between trial types was affected by condition. The trial type by condition mixed-design ANOVA revealed a main effect of trial type, $F(1,110) = 11.21$; $p < .001$, and a marginal main effect of condition, $F(1,110) = 3.55$; $p = .06$, both driven by the highly significant interaction, $F(1,110) = 26.72$; $p < .001$, $\eta^2 = .20$. To be more specific, trial types yielded different performance when participants made their predictions using real jars (jars-last condition: $M_{Big/heavy} = .88$, $M_{Small/light} = 0.67$; $F(1,110) = 36.94$; $p < .001$, $\eta^2 = .25$), but not when they made their predictions using pictures (jars-first condition: $M_{Big/heavy} = .78$; $M_{Small/light} = 0.83$; $F(1,110) = 1.63$; $p = .20$.)

Comparing post-test performance with training performance, we found no significant main effect of trial type in the jars-first condition, $F < 1$. Put differently, when adults were presented with pictures, they retained what they learned during the training and performed well even without feedback. In contrast, in the jars-last condition, when participants were given the opportunity to explore the objects haptically, performance changed from training to post-test. The 2 x 2 repeated-measure ANOVA (trial type by phase) revealed a significant main effect of trial type $F(1,56) = 14.83$; $p < .001$; a significant main effect of phase, $F(1,56) = 5.56$; $p = .02$; and a significant interaction, $F(1,56) = 42.63$; $p < .001$. From training to post-test, performance on big/heavy-wins pairs increased, $p < .001$, while performance on small/light-wins pairs decreased, $p < .001$. Put differently, participants in the jars-last condition reverted back to disambiguating the conflict in making predictions by focusing on the feature of weight.

Table 1. Summary of results.

Pre-test
• Independent of condition, performance was at ceiling when the winning jar was big and heavy.
• Independent of condition, systematic mistakes were made when the winning jar was small and light.
• Systematic mistakes were higher when participants used jars (compared to pictures).

Training
• Independent of condition, performance was equally high on both big/heavy-wins and small/light-wins pairs.
• Participants made some mistakes, but performance was overall above chance.

Post-test
• When participants used pictures, performance remained unchanged (compared to the training).
• When participants used jars, performance increased for the big/heavy-wins pairs, while it decreased for the small/light-wins pairs.

Difference scores. To capture these findings on the level of individual participants, we calculated a difference score for each participant, based on their performance on the big/heavy-wins and small/light-wins trials. Specifically, we subtracted average accuracy scores for the small/light-wins trials from the big/heavy-wins trials. This difference reflects the extent to which participants held a big/heavy bias, choosing the bigger/heavier jar as the winner more often than the smaller/lighter jar. Figure 6 shows obtained results.

Figure 6: Average difference between big/heavy-wins and
small/light-wins trials, per phase and condition.
Error bars represent standard errors of the mean.
Circles highlight when real-life jars were used.

Results are very much in line with our analysis of trial-based performance: Highest difference scores were obtained during pre-test ($M_{Pic} = .56$; $M_{Jar} = .71$), reflecting the naïve heaviness bias. Importantly, the difference score was higher for participants who were given the opportunity to explore the objects haptically than for participants who saw pictures, $F(1,110) = 4.74$; $p = .03$.

Difference scores decreased substantially during training, ($M_{Pic} = .02$; $M_{Jar} = .03$), reaching values that were statistically undistinguishable from zero, simple-sample $t < 1$. This suggests that participants no longer based their predictions on mass or volume alone. They quickly discovered a new criterion, yielding the same degree of success (i.e., number of mistakes) in both types of trials.

The central finding is during the post-test, after participants had learned about the shortcomings of a naïve heaviness bias. While participants in the jars-first condition largely retained the difference scores that they had obtained during training ($M_{Pic} = .04$), participants in the jars-last condition did not ($M_{Jar} = .21$). Their difference scores shot up, significantly more than the difference scores obtained in the jars-first condition, $F(1,110) = 26.72$; $p < .001$. Participants who were allowed to handle the real objects during the post-test phase seemed to unlearn some of what was learned during training.

Summary and Conclusion

We set out to explore the extent to which embodied experiences override, support, or interfere with experiences that are gained from visual perception. Adults participated in a prediction task about sinking objects, a task that is thought to elicit mistaken beliefs about what makes an object sink faster or slower in water. Feedback during part of the prediction task was expected to change some of those initial misconceptions. Indeed, adults demonstrated a substantial amount of learning during training. At the same time, we succeeded in creating a task that was sufficiently difficult for adults to perform below ceiling, but not so difficult that they would merely make random guesses. This is the kind of

regime that is likely to shed light on the constraints on mental activity.

How did embodied experience interface with performance when using 2D pictorial stimuli? Our results are clear: there was no evidence that embodied experience simply *overrode* visual perception. Even though adults were presented with the exact same trials across conditions, when their chance to explore objects haptically took place before training (jars-first condition), performance was decidedly different from when it took place after the training (jars-last condition). This suggests that behavior derived from embodied experiences is not separable from behavior derived from other means of perception. This, of course, is no surprise to a one-mind-one-behavior systems view (e.g., Clark, 2013; Smith, 2005). Visual and embodied perception are likely to be interlinked. Thus, results that argue for a dissociation between visual and embodied experience need to be re-evaluated carefully.

We also found no evidence that embodied experience *supports* visual perception. This is at least the case if support pertains to performing accurately. Whether participants got a chance to haptically explore objects before or after the training, their performance on the small/heavy-wins trials was lower with real jars than when they saw the objects as pictures. This is especially evident after the training, when participants reached equivalent levels of competence. Performance levels stayed the same during the post-test for adults presented with pictures, but critical mistakes arose from adults presented with the real-life objects. These findings undermine blanket claims of the general advantage of hands-on, embodied learning.

Results show that embodied experience *interfered* with visual perception. It did not act separately, and it did not support it, but nevertheless, it interfered with it drastically. This finding far from trivial given the current task, because relevant information, say about object mass and volume, were available to both modalities: participants could count the number of weights and compare the sizes of the objects, whether they were presented in real life or as pictures. If the same information can be obtained in theory, why then did we find differing performance as a function of condition?

Could it be that proprioceptive information simply made the task harder, yielding non-specific mistakes? This is unlikely, given that the differences in performance between the jar-based and picture-based contexts were rather specific, both in the pre-test and the post-test. In fact, there was not a general increase in mistakes for participants exposed to real-life objects: When they explored objects haptically, they performed highly successfully on big/heavy-wins trials, even better than adults who merely saw pictures. Their mistakes increased only on the small/light-wins trials. This pattern of performance, to perform well on big/heavy-wins trials and poorly on small/light-wins trials, is the signature of a heaviness bias, a bias that was more pronounced when participants could hold objects, rather than view them on a computer screen.

One could argue that our set-up was an unfair comparison: Embodied experience might support visual perception, but not in a task in which salience to heaviness yields mistakes. Embodiment might make heaviness salient, due to the inherently salient down-ward force of holding objects. Our results might reflect nothing more but a bias of embodied experience to increase the salience of heaviness, failing to generalize to embodied experience outside of heaviness tasks. While our data do not speak directly to this criticism, it is nevertheless worth questioning. This is because the difference in mass between the two objects in a pair is likely to be far more salient in the picture case than the real-jar case. The weights were too light (only 43g) to create differences that could be readily perceived haptically. It is most likely that adults judged difference in weight on the basis of visual information. Thus, a high salience of heaviness in embodied experiences might not explain our results.

There are several possible reasons for why the embodied experience increased the heaviness bias of adults. One possibility is that, rather than making heaviness more salient, the redundancy of information between visual and tactile information may have prompted the system to revert to a simpler belief (in this case, about heaviness). Without this redundancy of information, adults might have relied on their memory of feedback on specific pairs, and merely guessed on those pairs they could not remember. The haptic information might have disrupted this strategy to some extent. To test this possibility, the study would need to be expanded to include a manipulation of explicit, non-tactile disruption.

Our results support the idea that performance emerges from the interaction of many components that change each other over time driven by the system's own history (Smith, 2005; Smith & Breazeal, 2007). Such interactive processes have a non-linear character, and, beyond a certain size and number of relations among their constituents, they express a complex behavior of self-organization that cannot be explained by the simple features of the elements (Steenbeek & Van Geert, 2008; Van Orden, Holden, & Turvey, 2003). This dynamic pattern is difficult to place in a single component, because it is a product of the coordination of the whole system (Steenbeek & Van Geert, 2008). It is possible that the mind capitalizes on the dynamics of the body when needed (Spencer, Austin, & Schutte, 2012; Turvey, 1990; 2007). However, our results call for a more nuanced understanding of the effect of embodied experiences on the stability of representations.

Acknowledgments

References

Castillo, R. D., & Kloos, H. (2013). Can a flow-network approach shed light on children's problem solving? *Ecological Psychology, 25,* 281-292.

Castillo, R. D., Kloos, H. Richardson, M. J. & Waltzer, T. L. (2015). Beliefs as self-sustaining networks: Drawing parallels between networks of ecosystems and adults' predictions. *Frontiers in Psychology 6(1723).* doi: 10.3389/fpsyg.2015.01723

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(03), 181-204.

Chemero, A. (2011). *Radical embodied cognitive science.* MIT press.

Gibbs Jr, R. W. (2005). *Embodiment and cognitive science.* Cambridge University Press.

Goldin-Meadow, S., Cook, S. W., & Mitchell, Z. A. (2009). Gesturing gives children new ideas about math. *Psychological Science, 20,* 267-272.

Kloos, H., Fisher, A., & Van Orden, G. C. (2010). Situated naïve physics: Task constraints decide what children know about density. *Journal of Experimental Psychology: General, 139,* 625-637.

Kohn, A. S. (1993). Preschoolers' reasoning about density: Will it float? *Child Development, 64,* 1637–1650.

Smith, L. B. (2005). Cognition as a dynamic system: Principles from embodiment. *Developmental Review, 25*, 278-298. doi:10.1016/S1364-6613(03)00156-6

Smith, L. B., & Breazeal, C. (2007). The dynamic lift of developmental process. *Developmental Science, 10*, 61-68.

Steenbeek, H., & Van Geert, P. (2008). An empirical validation of a dynamic systems model of interaction: Do children of different sociometric statuses differ in their dyadic play? *Developmental science*, *11*, 253-281.

Spencer, J. P., Austin, A., & Schutte, A. R. (2012). Contributions of dynamic systems theory to cognitive development. *Cognitive Development*, *27*, 401-418.

Turvey, M. T. (1990). Coordination. *American Psychologist, 45*, 938-953.

Turvey, M. T. (2007). Action and perception at the level of synergies. *Human Movement Science, 26*, 657-697.

Van Orden, G. C., Holden, J. G., & Turvey, M. T. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology: General*, *132*, 331.

Wilson R. A. & Clark, A. (2009). How to situate cognition: Letting nature take its course. In Murat Aydede & P. Robbins (eds.), *The Cambridge handbook of situated cognition*. Cambridge University Press.