

Interactive Coin Addition: How Hands Can Help Us Think

Hansjörg Neth (neth@mpib-berlin.mpg.de)

Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development
Lentzeallee 94, 14195 Berlin, Germany

Stephen J. Payne (s.j.payne@bath.ac.uk)

Department of Computer Science, University of Bath
Bath BA2 7AY, United Kingdom

Abstract

Does using our hands help us to add the value of a set of coins? We test the benefits and costs of direct interaction with a mental arithmetic task in a computerized yoked design in which groups of participants vary in their interactive mode (move vs. look) and the initial configuration of coins (pseudo-random vs. another mover's final layout). By assessing performance and conducting a microgenetic analysis of the strategies employed we argue that the purpose of movement is the result, rather than the process of moving. Participants move coins in order to sort, rather than to mark, and select them by value, rather than by location. They spontaneously create remarkably smart solutions, thereby incidentally creating physical configurations that can help other problem solvers.

Keywords: embodied cognition, epistemic actions, complementary strategies, immediate interactive behavior.

Introduction

Do our hands help us to think? Although the roles of actions and gestures for guiding thought have often been recognized (e.g., Kirsh & Maglio, 1994; Clark, 1997; Cary & Carlson, 2001; Neth & Payne, 2001; Neth et al., 2007) the transformative potential of using our hands to re-organize and re-structure a problem is still poorly understood.

Consider the simple task of adding the value of coins scattered across a surface. Kirsh (1995) showed that a 'hands' condition was faster and more accurate at this task than was a 'look only' condition. However, the 'hands' condition was restricted to pointing to photographs of arrays of coins—moving the coins was not possible. Furthermore, beyond the important definition of *complementary strategies*—organizing activities that recruit external elements to reduce cognitive loads—Kirsh offers no precise account of the actual strategies which use of hands facilitated. In the absence of such an account it almost seems as if using hands provides a 'magic ingredient' for problem solving.

One rather mundane possibility is that the interactions *per se* have little effect, and all that matters is the information revealed in the resulting states of the world. This possibility has been supported with respect to interactive computer visualisations of an unfamiliar 3-D object by Keehner, Hegarty, Cohen, Khooshabeh, and Montello (2008). In these authors' experiments any learning benefit of interaction was matched by mere exposure to the most informative displays. In our experiment we test the force of a similar hypothesis with respect to moving coins while counting them.

A richer view of the role of external displays and their manipulation is provided in the work of Carlson and colleagues. Cary and Carlson (2001) showed that the availability of pencil and paper allowed participants to develop strategies more reflective of the structure of the task. Similarly, Carlson, Avraamides, Cary, and Strasberg (2007) showed that pointing increased both accuracy and speed in counting arrays of items and suggested that this was in part due to the provision of external markers for the boundaries between phases of a cognitive strategy. In this article we further explore this idea that movement allows the marking of items' roles.

We repeat Kirsh's (1995) experiment, with actual movement of coins allowed. Conducting our study in a virtual environment is still different from using real coins and multiple fingers, but allows us to address a range of subtle questions that go beyond the intuitive result that using hands can help us solve a mental addition task. Specifically, we answer the following questions:

1. *Performance benefits or trade-offs:* Does it help to use our hands to solve this arithmetic task? If so, is there a benefit of both accuracy and speed, or rather a trade-off, e.g., a benefit in accuracy with a corresponding cost in speed?
2. *Process vs. result:* Are the benefits based on the process of moving coins or the result of such a process? Are there systematic patterns in which objects are selected and moved?
3. *Purpose of interactions:* Are objects moved to mark those that are processed or to sort to facilitate future processing? Are movers reacting adaptively to the complexity of the task? Can use of hands also decrease performance by tempting humans to act instead of add on easy tasks?

To answer the performance-related questions we compare the tasks of adding sets of coins which can only be looked at with adding sets of coins which can be relocated during addition. Addressing the more detailed questions about *how* and *why* movement helps requires more methodological refinement. Rather than using physical coins, we present realistic life-size pictures of coins that can be moved on a computer screen by a mouse-operated drag-and-drop operation. This allows us to record complete behavioral protocols and perform detailed analyses of movement strategies.

Moreover, the use of a virtual environment enables us to incorporate a comparison between participants who receive specific spatial arrangements of coins and participants, who



Figure 1: Example screen display of an *initial* configuration (Participant 3, Trial 9).

are presented with somebody else’s initial or final coin configurations. This ‘yoked’ design allows us to address a related issue of *agency*: Is it important that people can interact themselves with the objects on display, or will they also benefit from the result of someone else’s interactions? To preview our results, we find that people successfully use physical movements of objects to transform the task and adaptively create structures that are helpful to other problem solvers.

Experiment

Method

Participants Sixty psychology undergraduates (44 female, 16 male) participated in the experiment in partial fulfillment of a course requirement. The mean age of participants was 20.6 years (ranging from 18 to 33).

Apparatus The experimental software was programmed in Microsoft Visual Basic™6 and ran on an Intel Pentium™4 PC with a 17” flat panel display.

All images of coins were distributed within a white rectangular display area, which measured 30 cm horizontally and 18 cm vertically at the 1024 by 768 screen resolution. When viewed from a 60 cm distance a coin was viewed at an 1.7–2.7° angle and the display area extended 26° horizontally and 16° vertically.

Materials Each stimulus consisted of a set of n British coins c_i , $c_i \in \{1 \text{ pence}, 2\text{p}, 5\text{p}, 10\text{p}, 20\text{p}, 50\text{p}, £1, £2\}$, with $n \geq 7$ on any trial. (See Figure 1 for an example.) 24 different stimuli sets (comprising a total of 396 coins) with a mean number of 20 coins (ranging from 12 to 21) and a mean value of 480 pence (ranging from 209 to 820 pence) were used.

Coins were presented as naturalistic photographs in their original dimensions, with diameters ranging from 18 mm for the smallest (5p) and 28 mm for the largest (£2) coins. However, as only a single image was used for each denomination, multiple coins of the same value were always viewed from the same side and at the same rotation angle.

Table 1: Overview of the four experimental groups.

		<i>interactive mode</i>	
		move	look
<i>configuration</i>	initial	1: <i>initial-move</i>	2: <i>initial-look</i>
	final	4: <i>final-move</i>	3: <i>final-look</i>

Note. Any differences along the *configuration* dimension depend on the actions performed in Group 1.

The stimulus sets also involved three dimensions that determined the complexity of the problem (by manipulating the number and values of coins, as well as the roundness of their sums). Due to space limitations, we do not analyze or discuss these within-subjects factors in this report.

Design A 2×2 variation of two between-subjects factors yielded the four experimental groups depicted in Table 1. The two between-subjects factors were:

- *Interactive mode*: In a *move*-condition participants were free to move coins individually by using a drag-and-drop procedure. In the *look*-condition participants could view the sets of coins, but could neither move or point at them. (To prevent its use as a pointing device the cursor was hidden within the display area.)
- *Configuration*: In the *initial* condition coins were scattered pseudo-randomly within the display area (with the constraint that two coins could not appear at the same location). The other’s final condition—hereafter *final*—depended on the manipulations of a particular person in the *initial-move*-condition and presented the configuration of coins at the end of the trial as the initial stimulus for another person.

For the sake of brevity, *members of* the four experimental groups will be referred to as initial movers (Group 1), initial lookers (2), final lookers (3), and final movers (4).

Note that the 2×2 variation of *interactive mode* and *configuration* separates the process of moving from the results of movement, and movers always have the option of *not* moving anything. Thus, any differences along the configural dimension depend entirely on the actions of movers in Group 1. In the extreme case of a particular mover deciding not to move anything all three yoked participants receive identical configurations of coins. Likewise, and with respect to the interactive dimension, any non-moving participants in a *move* condition would be self-selecting themselves into the corresponding *look* condition.

Group membership of participants to one of the four groups resulting from the 2×2 variation of the two between-subjects factors was assigned according to their order of arrival in the experimental laboratory. The presentation order of stimuli was randomized for every participant in Group 1. Yoked Groups 2 to 4 received the same stimulus order as the corresponding participant in Group 1.

Table 2: Addition accuracy: Mean number (and standard deviations) of errors by *configuration* and *interactive mode*.

<i>configuration</i>	<i>interactive mode</i>	
	move	look
initial	5.27 (3.56)	11.33 (6.49)
final	4.93 (3.99)	6.20 (3.65)

Procedure Participants were tested individually. They were instructed to add as quickly as possible without making errors until a criterion of 24 correct trials was reached.

Before the test trials participants practiced the procedure of adding items and entering sums on four lists of single-digit numbers and demonstrated knowledge of the drag-and-drop procedure in a scrabble task (when in the *move*-condition). Participants were also shown an example image of each coin and had to identify its value to the experimenter.

On a typical trial during the test phase, participants were presented with a stimulus, which they inspected or interacted with according to their interactive means. When they had completed adding, they pressed a key, whereupon the stimulus disappeared, and then were prompted for the result. Immediately after entering a result, participants received feedback regarding the correctness of their answer, and could request the next trial by pressing a key. This cycle was repeated until the participants had correctly added 24 stimulus sets. On average, participants completed the experiment within 32 minutes.

Hypotheses

The ability to interact with the stimulus set in both *move* conditions allows for a variety of functions, including pointing, marking, and sorting. Using some or all of these functions could allow movers to re-structure the mathematical properties of the task (e.g., exploit the commutativity and associativity of addition, or turn an addition of n identical values v into a multiplication $n \cdot v$).

As the strategies available to the participants in our *move* condition form a super-set of those available to participants in the *look* condition, our most basic hypothesis is that the former will exhibit performance benefits (in overall speed, accuracy, or both) over the latter.

If movers mainly move for moving's sake, i.e., benefit from the *process* of moving, then initial and final movers should show similar extents of movements, and receiving someone's final configurations is unlikely to convey a substantial benefit. However, if a main incentive for moving is the *result* of having moved objects then final movers should move objects to a lesser degree than initial movers and receiving the result of someone's moves is likely to convey a benefit.

We also hypothesize that two basic (not mutually exclusive) goals of movements could be to *mark* counted coins vs. to *sort* coins into clusters. Similarly, coins that are moved

Table 3: Addition latency: Mean trial times (and SDs) for correct additions by *configuration* and *interactive mode*.

<i>configuration</i>	<i>interactive mode</i>	
	move	look
initial	29.4 (8.0)	25.2 (6.1)
final	19.1 (6.0)	19.2 (6.1)

could be selected by their *location* or by their *value*. In the absence of any *a priori* assumptions we will sketch some descriptive analyses that inform these issues.

Results

We will first report basic performance results before considering the process of moving coins and the strategies, goals, and adaptivity of movers.

Performance

Accuracy Accuracy was analysed in terms of the total number of erroneous trials per participant. An ANOVA was performed with *configuration* and *interactive mode* as between-subjects factors. There was a significant main effect of both *configuration* and *interactive mode* and a significant interaction between *configuration* and *interactive mode*. The corresponding means and standard deviations are shown in Table 2. An analysis of simple main effects demonstrated that the combination of being shown the initial configuration and not being able to move the coins leads to significantly more errors than being able to move the initial configuration of coins, $F(1, 56) = 13.15, p = .001$, or than looking at the final configuration, $F(1, 56) = 9.41, p = .003$.

In short, initial lookers were less accurate than any of the other three groups. More specifically, being able to move coins or receiving the final coin configurations reliably increased the accuracy of participants' additions.

Latency As the integrity of the within-stimulus dimensions is only preserved when participants determine the correct value of a set of coins, latency was analysed in terms of time to complete correct trials.

Table 3 contains the mean addition latencies for correct additions in all four groups. There was a significant main effect of *configuration* ($p < .001$) but no main effect of *interactive mode* or interaction between these two factors. Thus, participants who encountered coins in their final configurations were reliably faster than those who encountered their initial configurations (19.1s vs. 27.3s) but being able to move coins did neither result in faster nor in slower performance than when only looking at them (24.3s vs. 22.2s, respectively).

Interactive Process and Strategies

Our performance results have shown that the option to move coins yields some benefits: Initial movers were more accurate

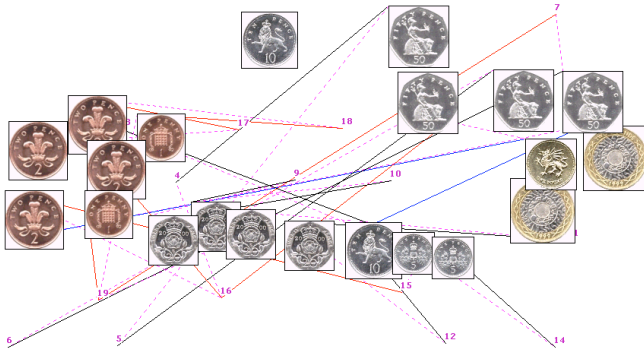


Figure 2: Example screen display of a *final* configuration (Participant 3, Trial 9). Thin solid lines indicate coin moves; dashed lines indicate distances between moves.

than initial lookers. However, so far we have only reported indirect evidence that any movement has actually taken place. The benefits (both in accuracy and in speed) derived from seeing the final configuration suggest that initial movers changed the display in a way that is helpful to final lookers.

In fact, participants in the *initial-move* and *final-move* groups made ample use of their option to move coins. Their recorded interactions comprise a total of 4,827 coin relocations, which were moved for a total distance of 457.8 meters in 2,953 seconds. (Figure 2 illustrates how an initial mover transformed the display of Figure 1.)

In the following, we will provide a precise process account of *what* initial movers did in order to explain the observed configuration effects. We also contrast the interactions of initial movers with those of final movers to assess the *purpose* and the *adaptivity* of their strategies.

Time allocations Table 3 reported total trial latencies as a function of group membership. But how did initial and final movers allocate their temporal resources? Table 4 shows their time distributions expressed as relative proportions of overall time. Although initial movers do not take reliably more time than final movers, Table 4a shows that they spend a vast majority of their total time (98.3%) on trials with moves (vs. 1.7% on trials without moves, $p < .001$). By contrast, final movers spend more equal proportions of their total time on trials with moves and on trials without moves (57% vs. 43%, $p = .301$).

Table 4b shows the relative time distributions for trials with at least one move. In both groups, the time interval from the trial onset to the selection of the first coin is very brief. Although final movers spend a larger proportion of a trial's duration before selecting the first coin than initial movers (18.0% vs. 8.1%, $p < .05$), the absolute first-coin selection latencies of both groups do not differ significantly (2.1s vs. 2.2s, $p = .89$). The brevity of this interval—which must include the initiation of movement, the choice of and cursor movement to the first target coin—excludes the possibility that much adding or planning can have taken place beforehand. This also implies that most subsequent decisions of

Table 5: Descriptive indices of the extent of coin relocations.

	<i>initial-move</i>	<i>final-move</i>
Percent of coins moved ^{***}	61.09 (17.93)	13.80 (11.78)
Number of moves ^{***}	265.13 (93.30)	56.67 (48.88)
Mean distance per move (cm) ^{***}	9.85 (0.86)	7.31 (1.38)
Mean duration per move (sec) [*]	0.64 (0.16)	0.53 (0.14)

Note. *** $p < .001$; * $p = .05$ in independent t-tests ($df = 28$).

whether and which coins to move next are made on the fly in a truly dynamic fashion. As the time interval between moves (i.e., between dropping one coin and selecting another) must contain the selection of and movement towards the next target, it is no surprise that both groups spend more time between moves than moving.

It is noteworthy, however, that both groups spend substantial periods of time *after* making their last move. Although initial movers spend a smaller proportion of their trial time after the last move than final movers (29.4% vs. 48.0%), their longer overall trial times mean that they spend more absolute time after the last move (8.4s vs. 5.2s, respectively, $p < .01$).

With respect to the interleaving of cognition and action this suggests that there is little planning prior to moving, but some mental addition takes place after the final coin movement.

Extent of movement Our performance and temporal process results suggest that initial movers moved more than final movers. As ‘moving more’ could mean many different things Table 5 shows the percentage of moved coins, the mean number of moves, as well as two measures of move distance and duration. Participants in the *initial-move* group not only made more moves, but a coin on average was moved for a further distance and for a longer time. Further, the fact that a coin was moved repeatedly in only 5.1% of all cases indicates that the process is highly successful at first shot.

If final movers can use the result of their yoked initial movers, there should be a negative correlation between the percentages of moved coins by initial and final movers. Indeed, the data shows an inverse relationship between the extent of movement in both groups, i.e., if a particular initial mover moved a lot the yoked final mover had less of an incentive to move, and vice versa (Pearson's $r = -.48$, $p < .05$). This indicates that those final movers whose initial movers had invested more effort were less inclined to move coins.

Purpose of moving Beyond showing performance differences due to the potential for movement and demonstrating *that* coins were actually moved a deeper explanation has to address questions about *how*, *where* and *why* coins were moved. With regards to the purpose of movement we considered two principal *a priori* hypotheses: Coins could be moved *to mark* already counted ones, or *to sort* coins into different arrangements. In the first case, the spatial position of a coin would be used as a primitive cognitive artifact to distinguish

Table 4: Time allocations in percent^a by group.

Group	(a) on all trials		(b) on trials with moves (m)			
	w/o moves	with moves	before m	during m	betw. m	after m
<i>initial-move</i>	1.7%	98.3%	8.1%	23.4%	39.1%	29.4%
<i>final-move</i>	43.0%	57.0%	18.0%	10.4%	23.6%	48.0%

^a To make both groups comparable, all latencies are expressed as relative proportions. The percentages are based on a mean total time of 704.7s ($SD=191.8s$) for the *initial-move* group and lower mean total time of 457.4s (144.4s) for the *final-move* group, $t(28) = 4.0, p < .001$.

between counted coins and coins yet to be counted. In the second case, the spatial arrangement of coins would be altered to facilitate their subsequent addition.

The performance benefits in the *final-look* condition and the inverse relation between initial and final movers already favors the sorting over the marking hypothesis. To demonstrate that coins were really sorted by their value we conducted two analyses that compared the number of coin clusters in the initial and final configurations. Two kinds of clusters were distinguished: If coins are sorted into clusters of *identical* values the number of identical nearest neighbors (INNs, e.g., a 10p coin being closest to another 10p coin) should increase. Similarly, if coins are sorted into groups with a round sum of values the number of *round* value clusters should increase. The presence of a ‘round value cluster’ was registered whenever a coin plus any subset of its 1–5 direct nearest neighbors added up to the next higher round number (i.e., 10p for coins $C_v, 1p \leq v \leq 5p$, and 100p for $10p \leq v \leq 50p$). (See Figure 2 for examples of both cluster types.)

Table 6a shows that the proportion of INNs increases substantially in the *initial-move* group, but not in the *final-move* group. Similarly, Table 6b shows a large increase in the number of round value clusters for initial but not final movers. Thus, both measures provide additional support that movers move coins to sort them into clusters, but note that the final result of their movements is still far from perfectly sorted.

Similar analyses, which we can only sketch here, show that participants select coins *by value*, rather than *by location*. More specifically, movers do not minimize the distance between subsequent pick-up or drop-off positions of coins,

but first select high-value coins and then coins of identical or lower value in a step-wise fashion. Figure 3 illustrates this characteristic pattern for the trial of which the initial and final configurations were shown by Figures 1 and 2, respectively. To implement such a selection-by-value strategy, both eyes and fingers of movers must traverse enormous distances.

Discussion

Our results demonstrate that a potential for direct interaction has multi-faceted consequences for the performance and process of a cognitive task.

With respect to performance, being able to move coins increases the accuracy of additions without a corresponding cost in latency. However, seeing the result of someone else’s moves increases both accuracy and speed.

Coins were moved to sort them into clusters of identical and round values, rather than to mark them by location. They were selected by value, rather than by location. This suggests that aspects of the mathematical task (concerning the values of addends) guided and governed the physical events (on the spatial and perceptual-motor level) in a top-down fashion.

The spontaneity and swiftness of interactions suggests that the complex interleaving of eye- and hand-movements with mental processes required little planning or reflective thought. An inverse relationship between the amount of movements by initial and their yoked final movers shows not only a high degree of adaptiveness in their actions, but also that there is some intersubjective agreement about what constitutes a ‘good’ configuration.

We oversimplified by stating the difference between moving for moving’s sake (or for the *process* of moving) and moving to achieve a final configuration (or to obtain a *result*) as a strict dichotomy. Although both aspects are not mutually exclusive, we accumulated much converging evidence that—for this task—the result mattered more than the process.

Conclusion

We began this paper with the hunch that using hands would help to add the values of a set of coins. Although our results confirm this intuition, our analysis uncovered many unanticipated aspects of this task. For instance, the result that the ability to move increases the accuracy but not the speed of additions is in contrast to Kirsh (1995)’s results, who reported an increase both in accuracy and speed when using fingers to point at photographs of coin displays. However, it is notable




Figure 3: Plotting selected coin values by elapsed time shows a characteristic downward step function. (See Figures 1 and 2 for the initial and final configurations of the same trial.)

Table 6: Proportion of identical and round value clusters at the beginning and end of all trials by *group*.

<i>Group</i>	(a) Identical nearest neighbors		(b) Coins in round value clusters	
	initially	finally	initially	finally
<i>initial-move</i>	19.5% (1.4%)	46.9% (11.5%)	23.5% (1.7%)	50.1% (7.2%)
<i>final-move</i>	46.9% (11.5%)	49.1% (8.6%)	50.1% (7.2%)	52.2% (7.0%)

that our non-significant result regarding latencies also fails to support the plausible counter-hypothesis that physical interaction involves a trade-off between speed and accuracy, i.e., incurs a cost in time as a price for increasing accuracy.

Any fine-grained account of the interactions between cognitive and physical task components requires microgenetic studies like ours as a necessary first step towards building accurate cognitive process models of interactive cognition. Our finding that movers care more about the result of their movements than benefitting from the process *per se* demystifies the ‘magic’ of moving in a similar way as Keehner et al. (2008)’s study explained the benefits of multiple views of a 3D-object. Rather than offering some magic bullet that somehow facilitates tasks, the potential for interaction provides concrete means for re-structuring a task into more manageable steps (e.g., changing an addition of many different values into a multiplication of identical values). Interestingly, even participants with very poor adding skills (as indicated by their error rate on simple addition tasks) had no difficulty using their hands to exploit the mathematical properties of addition.

Overall, our study provides a detailed account of the interplay between organism, task, and task environment, and shows how embodied creatures recruit complex interactive resources in the service of a simple cognitive task. The result that movers move coins to sort them into clusters (rather than adding and marking counted coins) and that they select coins by value (rather than minimizing the distance between moves) shows that movers willingly traverse long distances with their eyes and fingers to facilitate a mental task. The increased accuracy of their solutions without a corresponding cost in time justifies their interactions as a smart investment.

The solution strategies that movers spontaneously select have the side-effect that others can benefit from the configurations that are created during problem solving. This effect was unintentional as our movers did not know that the result of their interactions would be shown to other people facing the same task. Such benefits derived from receiving someone else’s final state are similar to the phenomena of *indirect interactions*, e.g., through the trails left by other pedestrians (Helbing, Keltsch, & Molnar, 1997; Goldstone & Roberts, 2006) or through spontaneous pattern formation in social settings (Moussaid, Garnier, Theraulaz, & Helbing, 2009). As human cognition heavily relies on artifacts (like language and notational systems, Clark, 1997) we speculate that seeing someone else’s solution to a problem is a basic—albeit sometimes accidental—origin of culture and technology.

References

- Carlson, R. A., Avraamides, M. N., Cary, M., & Strasberg, S. (2007). What do the hands externalize in simple arithmetic? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(4), 747–756.
- Cary, M., & Carlson, R. A. (2001). Distributing working memory resources during problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(3), 836–848.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: The MIT Press.
- Goldstone, R. L., & Roberts, M. E. (2006). Self-organized trail systems in groups of humans. *Complexity*, 11, 43–50.
- Helbing, D., Keltsch, J., & Molnar, P. (1997). Modelling the evolution of human trail systems. *Nature*, 388, 47–50.
- Keehner, M., Hegarty, M., Cohen, C., Khooshabeh, P., & Montello, D. R. (2008). Spatial reasoning with external visualizations: What matters is what you see, not whether you interact. *Cognitive Science*, 32, 1099–1132.
- Kirsh, D. (1995). Complementary strategies: Why we use our hands when we think. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 212–217). Hillsdale, NJ: Lawrence Erlbaum.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4), 513–549.
- Moussaid, M., Garnier, S., Theraulaz, G., & Helbing, D. (2009). Collective information processing and pattern formation in swarms, flocks, and crowds. *Topics in Cognitive Science*, 1(3), 469–497.
- Neth, H., Carlson, R. A., Gray, W. D., Kirlik, A., Kirsh, D., & Payne, S. J. (2007). Immediate interactive behavior: A symposium on embodied and embedded cognition. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Meeting of the Cognitive Science Society* (pp. 33–34). Austin, TX: Cognitive Science Society.
- Neth, H., & Payne, S. J. (2001). Addition as interactive problem solving. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 698–703). Mahwah, NJ: Lawrence Erlbaum.