Distinguishing Levels of Grounding that Underlie Transfer of Learning

Lisa A. Byrge (lbyrge@indiana.edu)
Department of Psychology, 1101 E. 10th Street
Bloomington, IN 47405 USA

Robert L. Goldstone (rgoldsto@indiana.edu)
Department of Psychology, 1101 E. 10th Street
Bloomington, IN 47405 USA

Abstract

We find that transfer of learning from a perceptually concrete simulation to an isomorphic but superficially dissimilar text-based problem is sensitive to the congruence between the force dynamics common to both systems and the kinesthetic schema induced via action in the first, perceptually concrete, simulation. Counterintuitively, *incompatibility* between the force dynamics and the kinesthetic schema has a beneficial effect on transfer, relative to compatibility as well as an unrelated control. We suggest that this incompatibility between action and system dynamics may make the system's relational structure more salient, leading to a more flexible conceptualization that ultimately benefits transfer. In addition, we suggest that too much "action concreteness" in hands-on learning may actually limit transfer, by fostering an understanding that is tied to that action and therefore less available for transfer in situations where that action is no longer relevant.

Keywords: Transfer; Action; Education; Embodiment; Analogy.

Introduction

Proliferating evidence from embodied cognitive science indicates that people regularly transfer understanding from visceral, body-based, experience to more "abstract" domains, such as from physical movement to conceptualization of time (Boroditsky & Ramscar, 2002) and from zygomaticus (smiling) muscle contraction to assessment of humor (Strack, Martin, & Stepper, 1988). A different kind of transfer of understanding – as traditionally conceptualized – is of crucial importance in education. Much of our educational system is based on the hope that prior learning can be transferred to novel situations. Unfortunately, research has shown that people often have great difficulties in transferring knowledge to new contexts (Gick & Holyoak, 1980). Much of the literature on transfer has focused on transfer of abstract problem-solving skills, such as applying a previously-learned mathematical equation to a novel type of problem. However, recent work has demonstrated an implicit form of transfer from a simulated concrete physical system to a superficially dissimilar textual problem with an analogous goal (Day & Goldstone, in press). In this work, the perceptual and spatial concreteness of the simulation, along with its dynamic nature, are thought to facilitate the development of a mental model that can then be deployed for interpreting and solving the less intuitive transfer task. This transfer was demonstrated for simulations controlled with a neutral action (a mouse click) or with no action (passive observation), suggesting that the transfer occurs at the level of the force dynamics that both systems have in common.

However, the above finding of implicit, spontaneous transfer from physical experience raises an interesting question. Would some form of "action concreteness" in the training simulation – parallel to the perceptual and spatial concreteness found to be beneficial – enhance this transfer still further? Adding action that is compatible with the force dynamics of the system might foster a more powerful representation of the force dynamics, by inducing a "kinesthetic schema" that echoes the conceptual force dynamics of the system. Such an enhanced, visceral understanding may then be more readily accessed and transferred. There is a growing literature on "action compatibility" effects in domains ranging from sentence comprehension (Glenberg & Kaschak, 2002) to insight problem solving (Thomas & Lleras, 2009; Catrambone et al, 2006), documenting facilitation when there is congruence between an action and a more abstract but related process. The idea that action compatibility is beneficial is also prominent in education, where researchers often assume that for difficult concepts, increased physical support leads to better learning and transfer.

However, work on contextualization suggests an alternative perspective -- too much concreteness can have a detrimental effect on transfer, causing what is learned to be "bound" to the learning context and less able to applied in novel situations (Goldstone & Son, 2009). On this view, action compatibility may actually harm transfer. The purpose of the current studies is therefore to determine whether action compatibility acts as a scaffold for learning and transfer, or as an obstacle to generalization.

Experiments

To investigate the effect of action compatibility on transfer, we operationalized action compatibility as the congruence between the *conceptual force schema* of the system being learned and the *kinesthetic schema* induced by the actions used to control the simulation of the system. We used the training and transfer task from Day and Goldstone (in press), with which positive transfer was demonstrated when the goals – and therefore both solution procedures and conceptual force schemas – of the training and transfer tasks were the same.

In the current studies, we manipulate the compatibility of the kinesthetic schema with the conceptual force schema *during the training simulation*, while holding constant the compatibility between the conceptual force schemas in the training and transfer tasks. If the facilitative effects of action compatibility include supporting learning and transfer, performance on the transfer task in the Compatible condition should be superior to the Incompatible condition. Alternatively, if action compatibility acts to tie the learned knowledge to specific actions that are relevant for the training but not transfer scenario – thus interfering with generalization -- then the Compatible condition should perform worse on the transfer task.

Experiment 1

Participants 48 Indiana University undergraduates participated in this study for partial course credit. 3 participants were excluded from analysis for failure to follow the instructions.

Materials and Design The materials in this study were based on the materials used in Day and Goldstone (in press). The training simulation in that original study depicted a horizontally oscillating ball suspended between elastic bands, and a rightward-pointing fan that could supply a constant horizontal force. The simulation included neither gravity nor friction, so the system yielded perpetual oscillatory motion of the ball when the fan was not being used. Participants clicked on the fan to get the ball to a particular position, such as "maximizing" the ball's position by getting it to the far endpoint. Consistently applying the fan only when the ball is moving to the *right* resulted in gradually *increasing* the amplitude of the ball's oscillation (because the rightward fan force augments, or *boosts*, the rightward acceleration of the ball). Consistently applying the fan only when the ball is moving to the *left* resulted in gradually *decreasing* amplitude of the ball's oscillation (because the rightward fan force cancels out, or *opposes*, the leftward acceleration of the ball). Therefore applying the fan all the time resulted in no net change in the ball's amplitude.

In the current studies, we added a "slider" controller for the fan that is activated using a pronounced hand movement to the right or to the left, depending on condition. Successfully "maximizing" the ball requires applying the fan only when the ball moves rightward, which gradually increases the amplitude of the ball's oscillation until it hits the far endpoint. Therefore in the condition with the rightward slider ("Compatible condition"; Fig. 1a), the rightward movement required to activate the fan is coupled to the rightward movement of the ball. This coupling between the ball and hand provides the kinesthetic experience of "moving in sync." With the added visual information of the ball moving further and further rightward, participants report that their experience is of "boosting forces." Conversely, in the condition with the leftward slider ("Incompatible condition," Fig. 1b), the leftward hand movement for activating the fan is exactly *opposite* the direction of the ball and fan. This results in the kinesthetic experience of "opposing forces." In both conditions, the *conceptual* force schema for the system is the same -- the fan boosts the ball. In the Compatible condition, the conceptual and kinesthetic force schemas are congruent -- the hand boosts the ball and the fan boosts the ball. But in the Incompatible condition, the conceptual and kinesthetic force schemas are incongruent -- the hand opposes the ball while the fan boosts the ball.

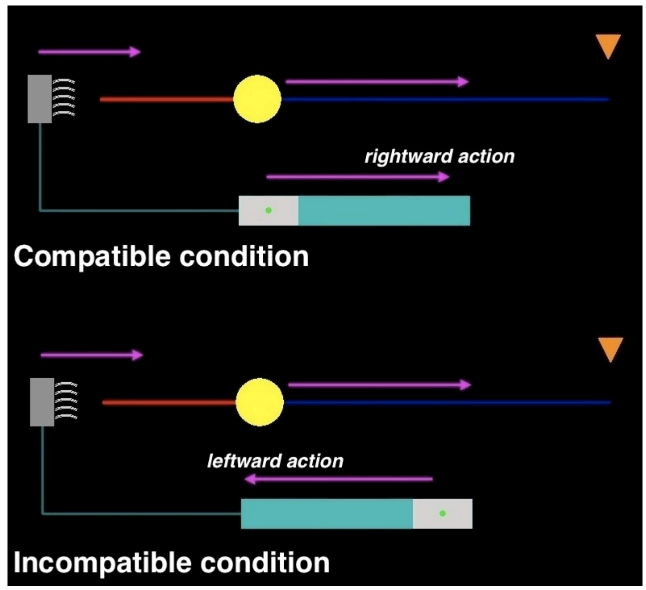

Figure 1a & 1b: Compatible (top) and Incompatible (bottom) conditions.
Video of these simulations is available at
http://cognitrn.psych.indiana.edu/complexsims/slider/max/

During the training phase, participants were guided through the ball simulation for six timed blocks. In Block 1, participants were asked to simply watch the system, with no interaction. With no interaction, there was no activation of the fan, so the ball simply oscillated regularly until the block ended. In Block 2, participants were asked to activate the fan and observe the differences between the ball's motion with and without the fan. Condition-specific instructions on how to activate the fan by controlling the slider with either rightward or leftward motion were provided (and repeated for every subsequent block). During Blocks 3-6, participants were instructed to observe the effect of using the fan in a specific way: "Try activating the fan ONLY WHEN THE BALL IS MOVING RIGHTWARD (away from the fan), and NOT when the ball is moving to the left." These instructions correspond to the correct solution procedure for getting the ball to reach the far endpoint ("maximizing" the ball). A participant who perfectly followed the instructions would therefore observe the "maximizing" of the ball's position four times. During all six blocks, the participants' task was to follow the instructions and observe the outcome; they were not asked to achieve any goal for ball's position. Although the far endpoint was highlighted with an orange triangle, participants were given no instructions pertaining to "maximizing" the ball or moving it to an endpoint. After completing Block 6, participants were instructed to ask the experimenter to start the second, ostensibly unrelated, experiment.

This second task was a strictly textual population dynamics problem identical to the transfer task used by Day and Goldstone (Fig. 2). Participants were presented with instructions describing a city that can comfortably hold 500,000 people; with less people, the city is more appealing due to low congestion; with more people, the city is less appealing due to crowding. All participants were given the task of getting the population to reach 1,000,000. Participants controlled the task by choosing whether to add media advertising, which increased the "appeal" of the city for one time step. This task proceeded in discrete time steps; at each time step, participants decided whether to click the "add media" button or the "no media" button. The value of the population, the appeal, and the change in appeal at each time step were displayed in a scrolling format, so that participants could also see the values for the previous five time steps. Participants could re-read the instructions at any time, and the goal ("get the population above a million") remained on-screen for the entire task. Participants were required to solve the task 3 times before completing the experiment.

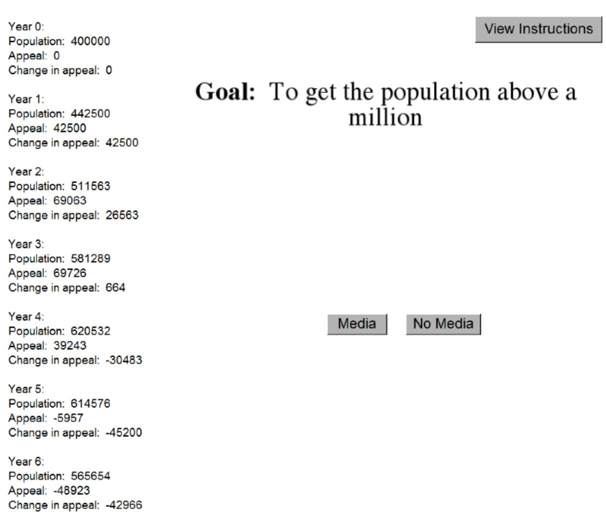

Figure 2: Population task.

Though seemingly unrelated in both subject matter and appearance, the two tasks are governed by the same dynamics: "population" is analogous to ball position, "appeal" to ball velocity, and "media advertising" to the fan. The behavior of both systems was thus governed by the same equations, with the population's value (or ball's position) oscillating around a numerical (or spatial) "midpoint." Just like turning on the fan in the ball simulation when the ball was moving rightward, adding media advertisement during the rise of the population increased the amplitude of the population's oscillation; adding media advertising during the fall of the population dampened the population's oscillation; and adding media advertising all the time caused no net change in the population's oscillation. Therefore, procedures for "maximizing" were also the same for the both systems, with participants intervening by adding media advertising (or fan force) to boost the population (or ball's position) only when it is already heading in the direction of the maximum.

The dependent variable in this study is the number of time steps needed to solve the population task on each of the three trials. We predicted that the different kinesthetic force schemas in the ball simulation would differentially affect transfer to the population task. If action compatibility facilitates transfer, the population task performance of the Compatible condition will be superior to that of the Incompatible condition. If action compatibility instead acts to bind what is learned to the specific actions, then population task performance of the Compatible condition will be worse than the Incompatible condition.

Results and discussion Most participants completed both simulations within the one-hour time limit. However, three participants from the Incompatible condition who failed to complete the population task within the time allotted were excluded from subsequent analysis.

There was no difference between the Incompatible and Compatible conditions in the average time steps required to complete the population task: 245.95 vs. 314.06 time steps ($t(327)=1.14$, $p=0.26$). However, a 3 × 2 (trial number × condition) mixed ANOVA revealed not only a significant main effect of trial ($F(2, 74)=5.30$, $p=0.0070$) representing improvement over time, but a significant interaction between trial and condition ($F(2,74)=4.156$, $p=0.019$). The interaction reflects an increasing advantage of the Incompatible condition over the Compatible condition as population trials increased from one to three (Fig. 3).

Multiple comparisons were conducted using Bonferroni-adjusted alpha levels of 0.0167 (0.5/3) per test. Results indicated that there was no reliable difference between the Incompatible and Compatible conditions on time to complete Trial 1 (456 vs. 330 time steps; $t(37)=1.055$, $p=0.30$, n.s.) or on Trial 2 (167 vs. 288 time steps; $t(37)=-1.87$, $p=0.070$, n.s.). However, the Incompatible condition significantly outperformed the Compatible condition on Trial 3: (117 vs. 324 time steps ($t(37)=-2.53$, $p=0.0157$).

Paired t-tests (Bonferroni-adjusted to alpha levels of 0.025) revealed that the time steps to complete the population task significantly decreased between the first and last trials of the population task in the Incompatible condition (a decrease of 339 time steps, $t(20)=3.40$, $p=0.002$), but not in the Compatible condition, which decreased by only 6 time steps. Because the Compatible condition's performance in the first trial was far from optimal – 330 time steps, while the shortest possible solution time is 39 time steps – the lack of improvement in that condition is striking.

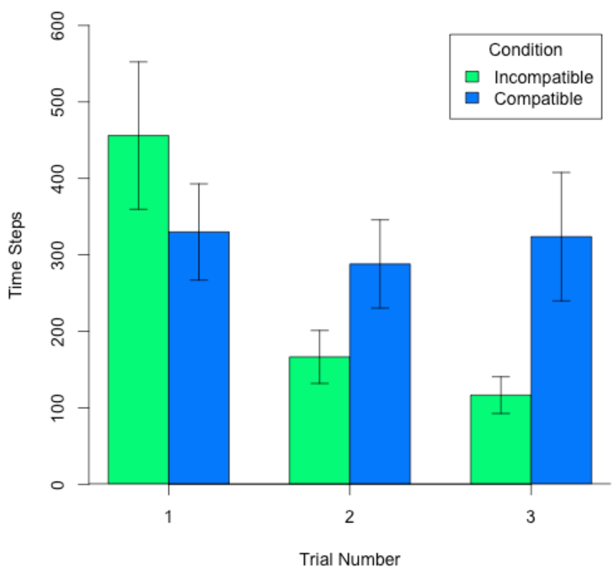

Figure 3: Population task performance. Participants who experienced the Incompatible training condition show more learning during an analogous population task than those in an action-compatible training condition. On this and all other graphs, error bars represent standard error of the mean.

Analysis of participants' strategy use in the population task provides converging evidence of a benefit for the Incompatible action training. We measured how often participants used the correct solution strategies, which meant adding media when the population is lower than 500,000 and increasing, and ceasing to add media immediately before (or, as soon as) the population was higher than 500,000 and decreasing. A 3 × 2 (trial number × condition) mixed ANOVA measuring total correct strategy usage revealed a significant main effect of trial ($F(2,74)=8.94$, $p=0.00033$) as well as a significant main effect of condition ($F(1,37)=4.83$, $p=0.034$) that reflected higher correct strategy use in the Incompatible condition (29% vs. 18%). Additionally, correct strategy use significantly increased between the first and last trials for the Incompatible condition (mean difference of 24%; $t(20)=3.24$, $p=0.0041$), but not for the Compatible condition (mean difference of 7.5%; $t(17)=1.49$, $p=0.15$, n.s.). A different strategy – adding media at every time step – was favored by both conditions initially, perhaps because it encapsulates the common-sense strategy of "if you want a large population, advertise your city as much as possible." A 3 × 2 (trial number × condition) mixed ANOVA revealed a significant main effect of condition ($F(1,37)=5.13$, $p=0.0295$) in usage of this intuitive, but incorrect, strategy, reflecting lower usage in the Incompatible condition (33% vs. 46%).

Finally, in the ball task, the Incompatible condition had a significantly lower proportion of trials where the "maximize" event was observed (0.51 vs. 0.75, $t(37)=-2.20$, $p=0.034$), indicating that the Incompatible condition had more difficulty in learning to move the slider according to the instructions. Apparently, the unintuitive actions required in the Incompatible condition slowed learning of the original ball scenario, but facilitated its transfer to the population task.

Experiment 2

Experiment 1 provides evidence that action compatibility in one task affects performance on an analogous transfer task. The relative transfer advantage of the Incompatible condition occurred in the absence of a problem-solving goal in the training task, showing that action compatibility matters even for non-goal-directed learning. However, the action-centered instructions in Experiment 1 may have led to a focus on the action itself, rather than the consequences of the actions. The positive transfer Day and Goldstone reported was primarily based on goal-directed learning of the ball task. Goal-directed learning is more likely to foster a functional perspective on the ball simulation, with a focus on the causal structure and the consequences of actions, rather than the actions themselves. If the action compatibility effect was a result of focus on the action at the expense of the system, then population task performance in the Compatible and Incompatible conditions should not differ after goal-directed learning of the ball simulation. However, if action compatibility genuinely affects conceptualization of the ball simulation, then differences should remain.

Additionally, Experiment 1 provides evidence of relative transfer differences between compatible and incompatible actions, but not that there is an absolute transfer advantage for either condition when compared to an unrelated training task. To address this, Experiment 2 also adds a control condition.

Participants 88 Indiana University undergraduates participated for partial course credit. 3 participants were excluded from analysis due to technical problems.

Materials and Design This experiment was nearly identical to Experiment 1, with the following important differences. A control condition was added, to assess whether the relative transfer advantage of the incompatible kinesthetic force schema was a genuine transfer benefit, relative to experience with an unrelated simulation. Participants in the control condition first performed a task involving repeatedly guiding a spacecraft to its home planet while being attracted to other fixed objects. We ensured that the spacecraft task required approximately the same time as the ball simulation. In the ball simulation, we added a goal: "maximizing" the ball. Participants received the same instructions on the operation of the slider as in Experiment 1, but the remaining procedural instructions were replaced with "Your goal in this task is: Make the ball reach the post opposite the fan, on the right side of the screen." To mitigate the effects of any initial difficulty in learning to control the Incompatible condition's slider, we also doubled the number of trials in which participants controlled the ball simulation. Therefore participants were required to solve the ball simulation 10 times before they were allowed to proceed to the population task. All participants, including the control condition, solved the same population task from Experiment 1.

Results and Discussion As in Experiment 1, most participants completed both simulations within the one-hour time limit. However, 5 participants who failed to complete the population task within the allotted time were excluded from subsequent analysis.

The Incompatible and Control conditions had significantly different average solution time steps in the population task, demonstrating a benefit for the Incompatible condition over the unrelated control (136.55 vs. 238.032; $t(51)=2.34$, $p=0.023$). A 3 × 3 (trial number × condition) mixed ANOVA revealed a reliable main effect of trial number ($F(2,154)=8.79$, $p<0.001$), with an effect of condition that approaches significance ($F(2,77)=2.69$, $p=0.075$) (Fig. 4). The population task performance exhibits the same trend as seen in Experiment 1, of an advantage of the Incompatible condition relative to the Compatible condition that emerges only after the first trial, although here the differences between those two conditions did not reach statistical significance.

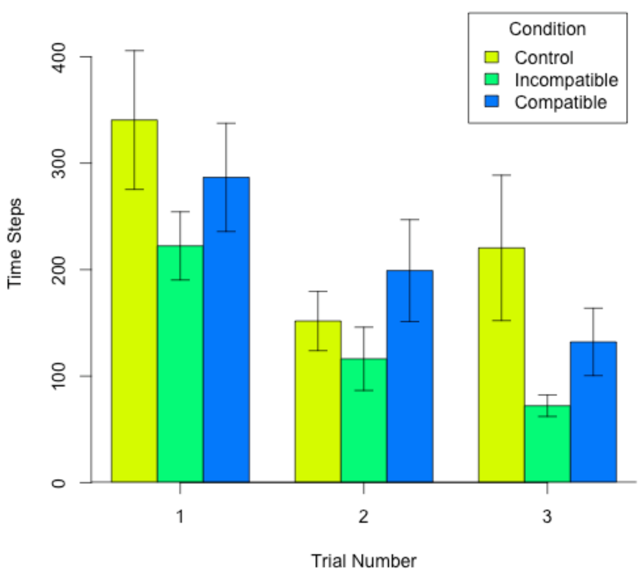

Figure 4: Population task performance following a version of the ball simulation or a control task, both with goals. The Incompatible condition results in better population performance than the control.

Comparison of population task strategy usage between the Incompatible and Compatible conditions echoes the strategy differences found in Experiment 1. A 3 × 2 (trial number × condition) mixed ANOVA for usage of one correct strategy -- adding media when the population is lower than 500,000 and increasing, and ceasing to add media immediately before the population was higher than 500,000 and decreasing (only this strategy is reported here because use of the other correct strategy did not differ) -- revealed a reliable main effect of condition ($F(1,47)=4.51$, $p=0.039$), demonstrating higher usage in the Incompatible condition (9% vs. 5%).

General Discussion

These two studies provide evidence of action compatibility differentially influencing transfer from a simulated physical system to an analogous but dissimilar task. The relatively better performance and strategy use in the Incompatible condition, compared to the Compatible or control conditions, support the notion that compatible actions are closely tied to what is learned such that when the functional force schema reappears in a transfer situation with different actions, there is a failure to transfer. In other words, action compatibility can act as an *obstacle* to generalization. Additionally, the superior transfer performance of the Incompatible condition relative to an unrelated control training task establishes that the incompatibility of the kinesthetic force schema and conceptual force schema offers genuine benefits for solving a transfer task with an analogous force schema.

Interestingly, despite worse performance on the *population* task, participants in the Compatible condition demonstrated better learning on the *ball* simulation in Experiment 1, as measured by their higher rate of successfully controlling the slider according to the instructions. This discrepancy between what participants find easy to do and what actually helps them is suggestive of the literature on "desirable difficulties." The conditions that facilitate immediate learning do not necessarily promote long-term learning (Bjork, 1994). On this view, the relative difficulty of controlling the Incompatible version of the ball simulation may have prompted deeper cognitive processing of the task and of the system as a whole.

While we believe that "desirable difficulties" contributes to the benefit of the Incompatible condition, it also leaves one aspect of this benefit unpredicted. If participants in the Incompatible condition simply acquired deeper, more transfer-relevant knowledge during the training task, then we might have expected *initial* differences between the conditions during the first transfer trials. And we might expect those initial differences to get *smaller* as time progresses, because the participants in the other conditions have had time to figure out a solution. Instead, we find the opposite pattern. All participants perform similarly initially, with differences appearing after the first trial.

The pattern of performance across trials can be interpreted under the framework of "preparation for future learning" (Bransford & Schwartz, 1999). On this view, the Incompatible ball simulation better *prepares* participants to *learn to do* the population task. Our prediction was in terms of a transfer disadvantage due to the unhelpful tying of learning to intuitive actions for the Compatible condition, but it is also instructive to ask what aspects of the Incompatible condition may be beneficial in itself.

One answer, we believe, is suggested by “conservative induction” (Medin & Ross, 1989). On this view, generalization emerges as a side effect of the use of specific examples, and the resulting generalized representation retains details of those specific examples. In the Compatible training condition, fan direction and hand motion were coupled together, controlling the ball that also moved in that same direction. It is conceivable, then, that the representation acquired in that condition was something like “things moving in sync.” Generalization to the population task may have then had the character of “acting in sync” with the oscillating population, or “apply media all the time.” In fact, this is precisely the “intuitive” and incorrect strategy the Compatible condition tended to use persistently throughout the population task. On the other hand, in the Incompatible training condition, the coupling between action and system dynamics is broken. The hand moves left while the fan is blowing right, which has the effect of moving the ball to the right. This decoupled version of the simulation was more difficult to learn to control. However, it may also have resulted in the acquisition of a representation that was more flexible, a representation that differentiated between the effect of the action for *controlling* the fan and the effect of the action *of* the fan. And by “teasing apart” – rather than “blurring together” -- the motoric and system-level actions, this more diagnostically structured representation contained the elements truly useful for connecting with, and solving, the population task.

Achieving transfer is difficult, but transfer is crucial to our educational system. These studies contribute to our understanding of a promising new form of transfer from concrete simulations, and demonstrate that while action compatibility can stand in the way of an educator’s goals, action incompatibility may sometimes help, by effectively isolating the relational structure from specific actions.

Acknowledgments

We thank Sam Day, Drew Hendrickson, and Paulo Carvalho for their helpful comments. We also acknowledge the support of National Science Foundation REESE grant 0910218.

References

Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe & A. Shimamura (Eds.), *Metacognition: Knowing about knowing.* Cambridge, MA: MIT Press.

Boroditsky, L., & Ramscar, M. (2002). The Roles of Body and Mind in Abstract Thought. *Psychological Science*, 13(2), 185-189.

Bransford, J. D., & Schwartz, D. L. (1999). Rethinking Transfer: A Simple Proposal With Multiple Implications. *Review of Research in Education*, 24(1), 61-100.

Catrambone, R., Craig, D. L., & Nersessian, N. J. (2006). The role of perceptually represented structure in analogical problem solving. *Memory & Cognition*, 34(5), 1126-1132.

Day, S., & Goldstone, R. L. (in press). Analogical transfer from a simulated physical system. *Journal of Experimental Psychology: Learning, Memory and Cognition.*

Gick, M. L. & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306–355.

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, 9, 558-565.

Medin, D. L., & Ross, B. H. (1989). The Specific Character of Abstract Thought: Categorization, Problem Solving, and Induction. In R. S. Sternberg (Ed.), *Advances in the psychology of human intelligence*. *Vol. 5.* Hillsdale, NJ: Erlbaum.

Son, J. Y., & Goldstone, R. L. (2009). Contextualization in Perspective. *Cognition and Instruction*, 27(1), 51-89.

Strack, F., Martin, L. L., & Stepper, S. (1988). Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology*, 54, 768–777.

Thomas, L. E., & Lleras, A. (2009). Swinging into thought: directed movement guides insight in problem solving. *Psychonomic Bulletin & Review*, 16(4), 719-23.