

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

On the limits of dynamic imagination: A mental extrapolation task

Permalink

<https://escholarship.org/uc/item/1751c9gz>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

ISSN

1069-7977

Authors

Levillain, Florent
Bonatti, Luca

Publication Date

2010

Peer reviewed

On the limits of dynamic imagination: A mental extrapolation task

Florent Levillain (levillain@lutin-userlab.fr)

Laboratoire “Cognitions Humaine et Artificielle”
Université Paris 8 – 2 rue de la Liberté 93526 Saint-Denis France

Luca L. Bonatti (lucabonatti@mac.com)

ICREA and Universitat Pompeu Fabra
Carrier Roc Boronat, 138, 08018 Barcelona, Spain

Abstract

To mentally extrapolate the trajectory of a moving object which disappears from sight, it is possible to exploit two different sources of information. One source is the memory of the last visible movement of the object, and the other is its inferred movement through time. It is often assumed that these cues are integrated into dynamical analog mental representations. To investigate the nature of the mental representation of imagined movements, we used a new experimental paradigm for which a causality attribution task was combined with motion prediction task. Participants were instructed to imagine the trajectory of a moving object disappearing behind a screen while estimating the degree to which the movement was caused by another moving object.

We show that the predicted movement departs from a correct extrapolation based on accurate memory for velocity. Furthermore the mental representation of the physical and causal structure of the dynamical events did not appear to be as detailed as a theory of mental simulation would predict.

Keywords: mental imagery, prediction of motion, perception of causality.

Introduction

Correctly performing actions on moving objects typically requires a high level of accuracy. Tasks, such as hitting or catching a ball show that humans can accurately and consistently represent the timing of a visible moving object and anticipate its future positions (Regan, 1982).

However, when the stimulus is not visible, such as when it is temporarily occluded, it is not clear how precisely we can time a non visible movement and whether we possess an extrapolation mechanism that can time non visible displacements. Interception tasks are mostly driven by kinematic properties, whereas mental extrapolation may be more influenced by cognitive factors, particularly by how we represent the causal interactions of the objects within a scene.

In studies of mental imagery, it is putatively assumed that the mind builds analog representations that can be used to estimate possible outcomes of dynamical events (Johnson-Laird, 1983) or to reveal spatial properties of objects (Kosslyn, 1994). Similarly, dynamical analog representations may subserve the ability to represent the timing structure of an invisible dynamical event (Shepard & Cooper, 1982; Schwartz, 1999).

Additionally, dynamical analog representations could integrate variables related to the physical structure of the

environment. Results suggesting that humans are capable of recognizing physically correct object movements (Kaiser et al., 1992), along with findings showing that we can perceive high-level properties of these stimuli, such as their causal relations (Leslie, 1994), or agency status (Premack, 1990), support this possibility. Indeed, it has been claimed that internalizing invariant properties of the environment is evolutionarily adaptive (Hubbard, 1995; Shepard, 2001).

Thus, it is plausible to conjecture that information regarding the dynamic properties of a scene that we are capable of representing (for example, their causal relations, or the amount of physical forces acting upon an object) is integrated in a unique mental simulation. This being the case, such a dynamical representation may allow for accurate prediction of future states of invisible events. Alternatively, the prediction of motion and the representation of other forms of physical information may be independent, and hence not merged into a single optimal simulation of dynamical events. In the present article, we aim to determine the ability to accurately estimate motions of invisible objects and to clarify how participants integrate an intuitive causal understanding of the represented events into a mental representation of motion.

Experiment 1

Experiment one determined the accuracy for predicting the position of a moving object that is no longer visible. Participants were required to predict the time-to-arrival of an animated ball at different positions after its disappearance.

We also tested how the representation of causal relations influenced participants' accuracy for predicting invisible dynamical events. If the information used to compute the velocity of an object is integrated with the information used to compute the causal structure of the scene, we would expect that events considered as causally correct are predicted more precisely than events considered as causally anomalous. However, if the two kinds of information are processed separately, we should observe a dissociation between the accuracy of online predictions of imagined position and the perception of causal correctness.

In every experimental condition, there were two moving objects, a launcher and a target, the movement of the target behind the occluder was to be predicted, while the causal relation between the launcher and the target, which could vary both in spatial and temporal contiguity, was to be estimated.

Method

Participants. Nineteen randomly chosen participants completed the experiment (mean age = 24,4; range from 20 to 31 years).

Stimuli. We created video stimuli with the animation software Cinema4D. The animation clips used were created with some intent of realism. For example, objects' shadows cast on the ground, had slight grooves, offering some depth cues. In each clip, a white and a green ball moved onto an earth-ground, below a blue, cloudy sky. A red screen partially covered the movement of the green ball (see Figure 1).

After 1 s, a white ball with a 3° diameter appeared from one side of the scene, and travelled horizontally at a constant speed of either 25,8°/s, 19,3°/s or 12,9°/s toward a green ball, which was stationary at the centre of the scene.

The white ball (the launcher) either did or did not contact the green ball (the target), but the target always started its movement as fast as the launcher and in the same direction.

A red rectangular screen was positioned such that its border contacted the edge of the target and its length covered the entire trajectory of the target. After initiating its movement, the target continued its trajectory behind the screen, until the end of the animation segment. Three vertical black lines were drawn on the red screen, placed at six different positions, yielding two configurations (see Tables 1 and 2). The direction of the balls (movement to the right, or left) was balanced across trials.

Three different spatio-temporal conditions were implemented by either varying the spatial interval between the launcher and the target at the end of the launcher's movement, or by varying the delay between the end of the launchers' movement and the beginning of the target's movement.

Table 1: Angular speed and hypothetical arrival time (s) in bar configuration 1 (in parentheses, distance from the origin of each bar).

Bar Number	12.9°/s	19.3°/s	25.8°/s
1 (44.2°)	0.44	0.28	0.24
4 (60.8°)	1.72	1.16	0.88
6 (71.8°)	2.6	1.72	1.32

Table 2: Angular speed and hypothetical arrival time (s) in bar configuration 2 (in parentheses, distance from the origin of each bar).

Bar Number	12.9°/s	19.3°/s	25.8°/s
2 (49.7°)	0.88	0.56	0.44
3 (55.3°)	1.28	0.84	0.68
5 (66.3°)	2.16	1.4	1.12

In the **Contact** condition, the motion of the launcher immediately ceased after having contacted the target, and the target began to move immediately after contact with the

launcher. Neither the launcher nor the target exhibited deformation as a result of contact.

In the **Delay** condition, an interval was introduced at the moment the two balls made contact. The interval was 480 ms for the first condition and 640 ms for the second condition.

In the **Space** condition, although the end of the movement of the launcher and the beginning of the movement of the target were simultaneous, the launcher stopped its trajectory before contacting the target. The space between the endpoint of the launcher's path and the target's starting position was determined according to the delays previously specified: the distance between the two balls was equal to the distance the launcher would have covered during the interval specified in the Delay condition had it continued its movement (a distance of 100 pixels for the first condition and 130 pixels for the second condition).




Figure 1. Overall sequence of events in a trial and causal conditions (Contact, Space, Delay).

To reveal the impact of occlusion on the time-to-arrival estimation, we designed another set of video sequences for which the target remained visible. However, in these clips the bars remained in the same positions as in the above described segments, and had the same spatio-temporal properties described previously. Also, the velocity of the target was constant and equal to the velocity of the launcher. Finally, in order to break the monotony resulting from the horizontal movements of the launcher, we intermixed the experimental animations with distractor segments in which the launcher fell from above, landing in the same position that the launcher stopped in the experimental sequences. Finally, for all the animations, the launcher appeared either from the left of the screen and moved right, or vice versa. Thus, in total, 120 experimental animations were created (5 conditions of interest crossed with the other experimental factors: Contact/Delay (x2)/Space (x2), target visibility (x2), bar configuration 1/2, speed (x3), direction of movement (x2) and 40 distractor animations).

A graded scale (from 0 to 9) was employed to collect participants' causal judgments for each clip.

Apparatus. Stimulus animations were displayed and the data were collected using a PowerMac G4 running the GNU software package PsyScope X (<http://psy.cns.sissa.it>). The animations were projected on a 200x135 cm screen with an Epson EMP 8100 projector. Reaction times were recorded using a Newmicros Button Box. This response box, together with a mouse and a numerical keypad were placed on a table positioned in front of participants.

Procedure. Participants sat in a darkened room, 2.5 meters from the screen. From their position, they could easily press the button box.

Each session began with a practice trial (Contact condition), with the velocity of the balls always set to 19.3°/s. Participants were instructed to visually track the launcher and, after the target disappeared behind the occluder, to press the key on the button box each time they felt the target would reach a bar on the red screen. They were encouraged not to press the key only three times. They were also informed that the balls would move at constant velocity and that they had identical speed. This information could be used to predict the position of the second ball on the basis of the speed of the first ball. Participants were able to move their head freely as they tracked the balls.

Participants were also informed that at the end of each segment a 1-to-9 scale would be projected on the screen. They were instructed to evaluate the perceived strength of the causal relation between the two balls by moving the mouse on the scale and clicking on the appropriate magnitude (1= not at all causal; 9 = completely causal). No explicit relation was drawn between the first online prediction and the second causal judgment tasks.

Each trial was initiated by pressing a button on the response box. The movement of the launcher started one second after the beginning of the scene. According to the velocity of the balls, the trial could last either 10, 11 or 12 seconds. At the end of the trial, a black screen, in which the causality scale appeared, filled the scene. After participants punched a number on a numerical keyboard, the next trial started. The beginning of the novel segment was controlled by participants. No feedback about response accuracy was given. Animation segments were presented in blocks of 80 (60 experimental, 20 distractors), arranged in a semi-random order, with the constraint that the same spatio-temporal condition could not be presented more than three times in a row. The first block contained only animations with occluded targets, and the second block contained the sequences with visible targets. The overall duration of the experiment was one hour, with a pause between the two blocks after thirty minutes.

Results

The mean timing error was computed as the difference between the total response time to a tested position from the beginning of the sequence and the total arrival time of the target, from the beginning of the sequence to the moment the target crossed a bar. The frame in which the invisible target ball reached each bar was determined offline, as the first frame where the target made contact with the bar. Thus, a positive error value indicates that participants entered their

response after the target crossed the bar, while a negative error value indicates the response was given before the arrival time.




Figure 2: Mean timing error (in ms) for the main experimental conditions (Contact/Space/Delay), separated by target visibility. The horizontal axis indicates the position of the bars on the trajectory of the target.

Trials were excluded from analysis if participants did not press the button exactly three times, if they pressed the button before the disappearance of the target, or if reaction times exceeded 2.5 SD from the mean response time in the relevant conditions.

We initially analyzed how the visible and invisible conditions differ for each tested position. Figure 2 shows the time differences between participants' responses and arrival time of the target ball, plotting together the three tested positions of each experimental animation sequence (1,4,6 and 2,3,5). When the target was visible, participants were accurate at determining the exact moment of arrival. Not only does the result show participants' accuracy at predicting contacts with direct visual feedback, but it also reveals that the task of tracking three successive positions with the spatio-temporal parameters we tested is perfectly feasible. However, when the target was not visible, clearly the average prediction systematically overestimated the time of arrival of the target. A two-way repeated measures ANOVA with individual means of timing error as a dependent variable and object occlusion as independent variable (visible/non visible) reveals an effect of occlusion on timing error ($F_{1,18} = 78.28, p < 0.0001$). The error was positive at every tested position, indicating an overestimation of the time needed for the target to cover the distance between the different bars. This response delay increased with occlusion time, but neither linearly nor continuously, as a simulation hypothesis would predict. In fact, the timing error did not differ between any two close bar pairs tested in the two bar position conditions. In other words, participants could not distinguish between any two close positions (as confirmed by post-hoc t-tests with Bonferroni alpha adjustment) when tested separately, but only in the invisible target condition, suggesting that the ability to predict the position of an invisible target is, if at all present, rather coarse.

In successive analyses, we thus collapsed the two bar position conditions, which did not differ. A two-way repeated measures ANOVA, restricted to the invisible target condition, with the sequence of response as independent variable (R1/R2/R3) revealed an effect of response order ($F_{2,18} = 109.04, p < 0.0001$). For each response, the timing delay increased as confirmed by a post-hoc t-tests (Bonferroni alpha adjustment) on the differences between timing errors in the third, second, and first response (**R2 – R1** = 631.02, $p < 0.0001$; **R3 – R1** = 1348.62, $p < 0.0001$).

Effect of causal conditions on timing accuracy. Figure 2 shows the timing error at each tested bar for the main conditions of the experiment. Timing was not accurate in any of the three tested conditions compared with the error values obtained in the visible movement conditions. There were also differences observed within the three conditions in the target-occluded animations. A two-way repeated measures ANOVA performed on individual error means for the invisible target conditions, with the type of interaction between the balls as the independent variable, indicated a significant effect of the three causal conditions ($F_{2,18} = 84.43, p < 0.0001$). This main effect depended on the difference between the Contact and the Delay conditions (post-hoc, Bonferroni adjusted t-tests : **Delay - Contact** = 337.84, $p < 0.0001$) and between the Space and Delay conditions (post-hoc, Bonferroni adjusted t-tests : **Delay – Space** = 293.14, $p < 0.0001$), while there was no difference between the Contact and Space condition (post-hoc, Bonferroni adjusted t-tests : **Space – Contact** = 44.7, $p = 0.53$).

Thus, the Delay condition reveals a timing error that is not only greater than the (almost null) error in the corresponding visible condition, but also greater than both non-visible conditions. Instead, timing errors in the Contact and Space conditions remained relatively similar across the trajectory of the target.

The effect of the spatio-temporal conditions did not depend on the size of the temporal or spatial intervals, as indicated by the lack of interaction between the two factors in a two-way ANOVA restricted to the responses in the Space and Delay conditions ($F_{1,18} = 0.02, p = 0.89$).

Effect of causal conditions on causal attribution. We analyzed participants' estimates of the causal strength of the scenes. A two-way repeated measures ANOVA with causal attributions as the dependent variable and interval size and causal conditions as independent variables revealed no effect of interval size ($F_{1,18} = 3.39, p = 0.08$) and no interaction between size and conditions ($F_{1,18} = 0.66, p = 0.43$). Thus, we collapsed the data across the interval dimension in further analyses.

A two-way repeated measures ANOVA yielded a significant effect of the type of interval introduced ($F_{2,18} = 66.99, p < 0.0001$). The effect was mainly carried by the difference between the Contact condition and Space conditions, but all conditions were different (post-hoc, Bonferroni adjusted t-tests : **Contact - Space** = 6.23, $p < 0.0001$; **Contact - Delay** = 4.52, $p = 0.03$; **Delay - Space** = 2.18, $p < 0.01$).

As expected, in the Contact condition the relation between the launcher and target was considered to be the causally strongest. Instead in the Space condition causality was considered non existent. Noticeably, causal interaction in the Delay condition was judged higher than in the Space condition. Combining such results with the prediction task, and comparing the two conditions in which causal violations were introduced, one can see that participants were better at predicting the position of an invisible target in the condition (Space) that was judged causally weaker than the other (Time). That is, prediction abilities and perception of causality do not align.

Effect of expertise on timing accuracy. Because many of our participants were highly skilled in physics and had a thorough understanding of real kinematics, we also checked whether expertise had any effect on accuracy. We divided the total number of participants in three groups based on the number of years they received physics education (naive: up to middle school; intermediate: up to high school; high: Masters and Ph.D in Physics).

Overall, expertise had no effect on prediction accuracy, as revealed by a two-way repeated measures ANOVA with individual means of timing error as a dependent variable and levels of expertise as independent variable (naive/intermediate/high) ($F_{2,18} = 0.35, p = 0.71$). Nor did any effect appear when causal attributions were the dependent variable ($F_{2,18} = 1.35, p = 0.29$). Expertise did not interact with spatio-temporal conditions in either predictive accuracy or causal attributions.

Discussion

Experiment one revealed that participants were highly accurate when predicting the time of contact of a moving target when the target was continuously visible, regardless of the type of interaction with the launcher. Yet, they were highly inaccurate when the target moved behind an occluder, making errors as high as 70% of the duration of the full scene. Furthermore, the amount of overestimation did not appear to increase continuously as the distance of the arrival point increased, revealing a sort of quantization of the error that is difficult to reconcile with a simulation theory of imagined movement.

This overestimation is difficult to explain by the violations of causal interactions in the events presented, as a large overestimation error was also present when the events were causally correct (Contact condition). Although we cannot be certain that, at a perceptive level, the computation of causal interactions does not interfere with the prediction, we found that the attributions of causality were dissociated from prediction accuracy: participants were better at predicting the position of an unseen object in conditions that they judged causally worse.

It is thus more likely that the variations in the amplitude of timing error have a source in the time necessary to integrate the two successive movements at a purely kinematic level. As such, this experimental situation might reveal particularly interesting in the exploration of movements integration.

Overall, these results suggest that our ability to predict future states in a partially occluded dynamical event is severely limited and probably does not integrate our knowledge about causal interactions.

Experiment 2

An alternative explanation for the large delays observed in experiment one, which maintains the tenet that humans simulate physical events, could be that participants do simulate physical events, but they do it even better than required. We showed animation sequences in which balls rolled over flat terrain. If participants integrate real physical constraints they may not avoid considering friction in their simulation, thus 'mentally slowing down' the speed of an unseen object. This integration of a physical variable might explain why participants delayed their reactions in imagination. Although the size of the delays we found is not easily reconciled with a simple integration of real friction parameters given the terrain in our videos, the point remains valid. Indeed Hubbard (1995) suggested mental analogs of gravity and friction are directly integrated in our simulations of object motion, systematically biasing certain position estimations. So the time-to-arrival overestimation in our experiment could reflect the fact that participants are simulating a deceleration instead of using their memory of a constant velocity.

We tested this possibility by modifying the context of the previous sequences, so as to prime certain physical representations. Specifically, we tilted the slope of the track such that the balls would either roll downwards or upwards. A previous study has shown that such a transformation can bias memory for position in a representational momentum paradigm (Bertamini, 1993).

If the prediction is indeed driven by inferred dynamical properties, we expect the timing error to be modified according to the orientation of the slope. If, instead, the prediction is not affected by the integration of physical variables and the error we found in Experiment one was due to limits in how we can simulate physical events (if we have such an ability), then we expect the timing error to persist unaffected by the conditions of Experiment two.

Method

Participants. Thirteen randomly chosen participants were recruited for the experiment. Their ages ranged from 19 to 30 years (mean age = 22,9).

Stimuli. We used the animations with occluded target movement from Experiment one, but modified such that the slope of the track was altered by rotating the images 20° either clockwise or counterclockwise. Thus, in experiment two there were three groups of animation stimuli: two containing balls rolling on an inclined plane, and a third group containing the same sequences used in Experiment one, with the balls rolling on a horizontal plane. This configuration allowed us to determine how gravity modifies the results of Experiment one. No vertical movement distractor was present in this experiment.

Apparatus and Procedure. The same set-up and procedure used in Experiment one were used for Experiment two, with the exception that we did not run the visible target condition, as performance in this condition was previously shown to be accurate.

Results

Data exclusion criteria and error calculations were as in Experiment one.

Effect of the orientation of the slope on timing accuracy.

Figure 3 represents the variation of timing error as a function of slope. As in Experiment one, an overestimation of the time-to-arrival, increasing with response order, was observed. There was no obvious difference in timing accuracy between the different slope conditions, although a slight decrease in timing error appeared in the slope downward condition.

A two-ways repeated measures ANOVA performed on the timing error, with speed and slope as independent variables, revealed a main effect of slope ($F_{2, 12} = 4.61$, $p = 0.02$). Post-hoc t-tests with Bonferroni adjustment revealed that the effect was carried by the difference between the Slope up and Slope down conditions (**Slope up - Slope down** = 175.34, $p = 0.02$), whereas no difference was found between the two tilted conditions and the horizontal condition. Furthermore, the difference between these conditions only occurred at one speed. Indeed, speed and slope interacted ($F_{4, 12} = 2.87$, $p = 0.03$); post hoc analyses showed that the difference between Slope up and Slope down was significant only when the balls moved at 19°/s (post-hoc Bonferroni-adjusted t-tests: 19°/s **Slope up - Slope down** = 300,86, $p < 0.01$; 13°/s **Slope up - Slope down** = 102,38, $p = 0.99$; 26°/s **Slope up - Slope down** = 122,78, $p = 0.94$).




Figure 3. Mean timing error (ms) for the three plane rotations (Slope up, Slope Down, Horizontal). The horizontal axis indicates the position of the bars on the trajectory of the target.

Discussion

In this experiment, we tested whether the prediction of the position of an unseen object was influenced by the

integration of physical variables in a mental simulation of the dynamic of the action.

We observed a slight effect of slope on participants' predictions, but only at one velocity. While the effect is compatible with the mental simulation of physical parameters, it remains mysterious as to why it should occur only in the 19°/s velocity condition. Thus, overall, it is difficult to entirely reconcile the results with the assumption that our mental models faithfully simulate the dynamics of object movements.

General Discussion

How does the cognitive system deal with incomplete information about the trajectory of a moving object? Research on mental imagery and on the prediction of motion frequently appeals to mental analog representations as a potential substrate for spatial computations and dynamics understanding. Here we provided evidence that this conception may not offer an adequate account of how we represent dynamic stimuli.

We devised a task for motion prediction that directly probed participants' ability to estimate the position of a moving object online, as opposed to other known paradigms of motion prediction which test memory for past positions rather than fast prediction of future positions (e.g., Hubbard, 1995). With this task, we demonstrated that estimations of time-to-arrival are inaccurate, with a large overestimation of the time necessary for the target to reach a position (confirming and expanding upon previous results obtained with different paradigms; e.g., Gilden Blake & Hurst, 1995). This result supports the claim that there is no predictive mechanism to estimate an object's position when it is occluded, when a direct visual evidence is lacking (Keane & Pylyshyn, 2006).

Furthermore, by coupling this task with causal strength judgment task, we showed that intuitive perceptions of causality do not integrate with online prediction of imagined object movements, casting further doubts on the existence of a representation that integrates physical variables into an analog simulation of objects and physical forces in the world. Finally, we showed that the system responsible for the overestimation error we revealed, takes into account very obvious physical properties, such as gravity, only haphazardly. This aspect of our results is difficult to reconcile with evolutionary accounts of cognition, according to which integration of gravity should be a prime candidate for a variable that evolutionary history may have embodied into a mental simulator.

How then can we account for the overestimation error we observed? Some studies suggest that when we track a moving object, our time perception for rapidly moving stimuli is lengthened as compared to static stationary stimuli (Brown, 1995; Kanai et al., 2006). Such a phenomenon could account in part for the present results, and as a consequence it could indicate that rather than extrapolating object position by means of an analog mental simulation of real physical forces, we use an internal clock to make an only coarse estimate of when an invisible object should be at a given location.

As a general conclusion, our results point toward the existence of several independent systems, one of which may compute object velocity, and another that may compute causal relations in the world. Although it may be tempting to unite the two kinds of systems, our results cast doubt on the existence of a common substrate for the extrapolation of trajectories in dynamical sequences of movements. These results also cast doubts on the existence of richly detailed analog representations that could assist us in knowing and understanding the physical world.

Acknowledgments

Our thanks to J. Mehler and the members of the Language, Cognition & Development lab in Trieste for helping us in designing and testing our experimental paradigm. This research was supported by a MICINN PSI2009-08232 grant to L.B.

References

- Bertamini, M. (1993). Memory for position and dynamic representations. *Memory & Cognition*, 21, 449-457.
- Brown, S. W. (1995). Time, change, and motion: The effects of stimulus movement on temporal perception. *Perception & Psychophysics*, 57, 105-116.
- Gilden, D., Blake, R., & Hurst, G. (1995). Neural adaptation of imaginary visual motion. *Cognitive Psychology*, 28, 1-16.
- Hubbard, T. L. (1995). Environmental invariants in the representation of motion: Implied dynamics and representation momentum, gravity, friction and centripetal force. *Psychonomic Bulletin & Review*, 2, 322-338.
- Johnson-Laird, P. N. (1963). *Mental Models: Towards a cognitive science of language, inference, and consciousness*. Cambridge University Press.
- Kaiser, M. K., Proffitt, D. R., & Anderson, K. A. (1985). Judgments of natural and anomalous trajectories in the presence and absence of motion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 795-803.
- Kanai, R., Paffen, C. L. E., Hogendoorn, H., & Verstraten, F. A. J. (2006). Time dilation in dynamic visual display. *Journal of Vision*, 6, 1421-1430.
- Keane, B., P., & Pylyshyn, Z., W. (2006). *Is motion extrapolation employed in multiple object tracking? Tracking as a low-level, non-predictive function*. *Cognitive Psychology*, 52(4), 346-368.
- Kosslyn, S.M. (1996). *Image and Brain: The resolution of the imagery debate*. MIT Press, London, England.
- Leslie, A. M. (1994). ToMM, ToBy, and agency: core architecture and domain specificity. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the Mind; Domain specificity in cognition and culture*. Cambridge University Press.
- Premack, D. (1990). The infant's theory of self-propelled objects. *Cognition*, 36, 1-16.
- Regan, D. (1992). Visual judgments and misjudgments in cricket and the art of flight. *Perception*, 21, 91-115.
- Schwartz, D. L. (1999). Physical imagery: Kinematic versus dynamic models. *Cognitive Psychology*, 38, 433-464.

- Shepard, R. N. (2001). Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences*, 24, 581-601.
- Shepard, R. N., & Cooper, L. (1982). *Mental Images and Their Transformations*. MIT Press.