

DISSERTATION

PERCEPTION OR RESPONSE BIAS? EVALUATING COMPETING HYPOTHESES THROUGH  
AUTOMATION OF ACTION-CONTROL

Submitted by

Nathan L. Tenhundfeld

Department of Psychology

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Fall 2017

Doctoral Committee

Advisor: Jessica Witt

Benjamin Clegg  
Alyssa Gibbons  
Charles A. P. Smith  
Christopher Wickens

Copyright by Nathan Laurence Tenhundfeld 2017

All Rights Reserved

## ABSTRACT

### PERCEPTION OR RESPONSE BIAS? EVALUATING COMPETING HYPOTHESES THROUGH AUTOMATION OF ACTION-CONTROL

The claim of action-specific researchers is that one's ability to act affects his or her perception of the environment. When using a reach extending tool, such as a stick, objects appear closer than they do without using that stick. However, whether these effects are perception or simply a response bias has been hotly contested. In this dissertation, four experiments were run using the Pong task to be able to differentiate between a response bias and evidence for a perceptual account. Results indicate that not only were results not in line with a response bias account, but they were what the action-specific account of perception would predict. Results are discussed in context of what this means for theories of visual perception. Results are then discussed in relation to the motor simulation hypothesis to evaluate its validity as an explanation for action-specific effects. Finally, given the nature of the experimental design, a framework for a Theory of Automation Embodiment is developed.

## ACKNOWLEDGMENTS

I would like to express my deepest appreciation to my Committee Chair, and Advisor, Dr. Jessica Witt. Without her guidance, patience, and flexibility, this dissertation, and the entirety of my graduate school training, would not have been possible. A better research mentor the world over cannot be found. I am still planning on coming to clean that coffee stain off your office carpet by the way...

I would also like to thank my committee members Dr. Ben Clegg, Dr. Alyssa Gibbons, Dr. CAP Smith, and Dr. Chris Wickens for their years of training, challenging, and pushing me. Your willingness to always entertain my most critical and stubborn side is uniquely appreciated. May every committee hereafter be required to have members so dauntingly brilliant and knowledgeable!

## DEDICATION

This dissertation is dedicated to my family and friends that provided uncompromising love and support throughout my time in graduate school. Including special recognition of “Grandad”, “Nene”, “Grandma”, and Aunt Lynn, each of whom were there with me at the start and throughout the ups and downs of this journey, but passed before I could cross the finish line.

Additionally, I want to specifically thank my parents, Mark and Lynda Tenhundfeld, as well as my brother, Danny Tenhundfeld, for their constant shared excitement in my victories, and unyielding support during my moments to learn from. I would not have survived graduate school without you three, and I love you each more than I will ever be adequately able to express. In life you always get to choose your best friends, but only the luckiest ones get to share a name with theirs.

“If you can fill the unforgiving minute  
With sixty seconds’ worth of distance run,  
Yours is the Earth and everything that’s in it,  
And – which is more – you’ll be a Man, my son.”

## TABLE OF CONTENTS

Abstract.....	ii
Acknowledgments .....	iii
Dedication .....	iv
Chapter 1: Action-Specific Perception .....	1
Chapter 2: Controversies Surrounding Action-Specific Perception .....	5
Chapter 3: Experiment 1 .....	11
Method.....	11
Results.....	15
Discussion.....	18
Chapter 4: Experiment 2 .....	22
Method.....	22
Results.....	26
Discussion.....	28
Chapter 5: Experiment 3 .....	30
Method.....	30
Results.....	30
Discussion.....	31
Chapter 6: Experiment 4 .....	34
Method.....	34
Results.....	35
Discussion.....	36
Chapter 7: General Discussion .....	39
Perception versus Response Bias.....	39
A Role for Motor Simulation .....	45
Framework for a Theory of Automation Embodiment.....	53
Applied Considerations.....	60
Conclusion.....	63
References .....	64
Appendix A .....	74

## Chapter 1: Action-Specific Perception

The action-specific account of perception suggests that our perceptual system embodies the capabilities of the body in order to formulate the percept of our surrounding environment (e.g. Witt, 2011-a). This account represents a significant departure from other theories of perception because it emphasizes action as a source of information for perception, rather than merely the end-goal. The research on the embodiment of the capabilities of the body can be broken down into three major categories: energetic effects, performance effects, and skill effects.

Energetic effects are generally classified by changes in energetic expenditure affecting one's perception. For example, individuals tend to report seeing distances on a hill as being farther than distances on the flat ground (e.g. Stefanucci, Proffitt, Banton, & Epstein, 2005; Tenhundfeld & Witt, 2017). The claim of embodiment is that because it takes more energy to traverse the same egocentric distance up a hill, when compared to the flat ground, the perceptual system is accounting for this change by distorting the percept to match this discrepancy in effort. A similar effect has been found in the energetic expenditures associated with throwing a ball. When tasked with throwing a heavy ball, participants reported seeing distances as farther than when throwing a lighter ball (Witt, Proffitt, & Epstein, 2004).

These energetic effects appear to be related to physical body size rather than beliefs about body size. Sugovic, Turk, & Witt (2016) found that obese individuals report seeing distances as farther than do their thinner counterparts. Interestingly, their results indicate that perception is influenced by actual body size (as measured by body weight), rather than

perceived body size (as measured by selection of an image that most closely matches one's own body). Again, this represents an example of an individual's ability to act in their environment changing how they perceive it.

The second broad category of these action-specific effects relates to changes in performance affecting changes in perception. Softball players who are hitting better report seeing the softball as larger than do those who are not hitting as well (Gray, 2013; Witt & Proffitt, 2005). Similar effects have been found in archery (Lee, Lee, Carello, & Turvey, 2012), putting (Witt, Linkenauger, Backdash, & Proffitt, 2008), and tennis (Witt & Sugovic, 2010) to name just a few. One interesting dilemma, however, is whether these softball players, archers, golfers, and tennis players are seeing things differently *because* of their performance, or, if their performance is better because of their perception.

In an effort to determine the directionality of these effects, one study surveyed participants before and after kicking field goals (Witt & Dorsch, 2009). Researchers found that there were no differences in the perceived size of the goalposts before kicking, but kicking performance predicted their post-performance perception of goalpost size. While it appears to be true that affecting perceptions of, for example, a golf hole can affect performance (Witt, Linkenauger, & Proffitt, 2012), these action-specific researchers have made the claim that the perceptual system is embodying one's performance, which distorts their perceptions accordingly.

Finally is the category of skill. Research has shown that free runners (i.e. those who do Parkour) perceive wall heights to be lower than do their free running novice counterparts (Taylor, Witt, & Sugovic, 2011). Because these free runners have developed the ability to scale walls better than your average individual, their perceptual system has

embodied this ability, which results in them perceiving the walls to be less tall. Similar effects have been observed with swimmers. Those who are better swimmers report seeing underwater distances as being closer than do individuals who are less skilled (Witt, Schuck, & Taylor, 2011). One explanation for these results is that the better swimmers are actually more energetically efficient, and thus their mind accounts for that in the formation of their percept.

These sort of reported perceptual distortions have also been found in direct manipulations of action capabilities. When reaching with a baton, objects that were otherwise out of reach but now can be reached look closer than they do when reaching without the baton (Witt, Proffitt, & Epstein, 2005; Witt, 2011-b). Similarly, when using flippers, swimmers see underwater targets as closer than when they are unaided by the flippers (Witt et al., 2011). Even cars seem to affect perception; when driving, distances look shorter than they do when walking (Moeller, Zoppke, & Frings, 2016). Each of these studies involved a tool that enhanced action capabilities, and found subsequent effects on an individual's perception. Thus, the effect of action on spatial perception can accommodate dynamic changes to the body via tool use.

However it is not simply the action itself that matters. Research has also expanded these effects into examinations of intention to act. As a definition of intention I have chosen to use one taken from Krueger (2003), which states that intention is “a cognitive state that is temporally prior and immediately proximate to the target behavior”. Therefore the intention to act could be seen as motor simulation of the anticipated action (Witt & Proffitt, 2008; Wraga, Thompson, Alpert, & Kosslyn, 2003). Research has shown that when asking participants to imagine using a reach extending tool, there is a similar perceptual effect to

as if the participant actually reached with the tool (Davoli, Brockmole, & Witt, 2012). Similarly, when participants anticipate being able to use the reach extending tool, even when they are not actually using it, there is a similar perceptual distortion to what would be expected if they did use that tool (Witt & Proffitt, 2008).

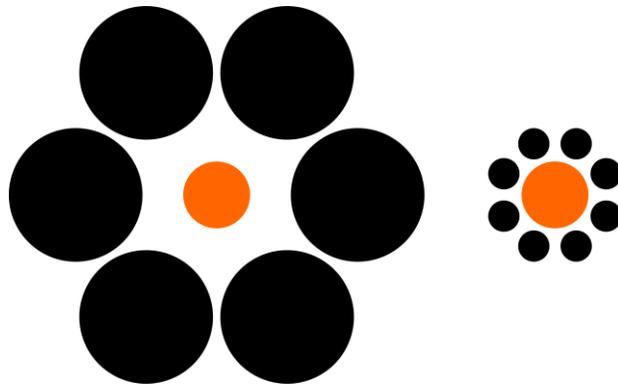
All of these aforementioned explorations into action's effect on perception have the same underlying theory, namely that a "perceiver's abilities are... reflected in perception." (Witt, 2011-a). What is remarkable is that in each of these action-specific perception studies, the athletes/participants have similar *optical* experience. Yet, differences in their perceptual experience were observed. The notion that the individual's ability affects how they perceive the world directly challenges the way many have viewed and discussed perception. Because of that, these results have been met with significant theoretical challenges.

## Chapter 2: Controversies Surrounding Action-Specific Perception

There has been substantial debate over whether action-specific effects actually represent a genuine effect on perception, or if they are simply due to response bias. Part of this debate stems from the fact that we cannot *directly* measure perception. We can only measure behavioral responses such as perceptual judgments or actions. These responses are driven by several processes in addition to perceptual processes, making it difficult to isolate the aspect of the response that is due to how the object was perceived versus due to the processes that help generate the response. We can develop tests to try and determine an individual's percept, but there is no specific way for researchers to directly observe what a participant sees. Because of this, alternative explanations for what Action-Specific researchers have called perception, abound. If the effects are nothing more than, for example, response biases, then they are not at all revolutionary. However, if it can be demonstrated that action genuinely has an effect on perception, this would demand current theories of vision to take these effects into consideration.

The notion that one's ability to act somehow directly affects perception of the environment stands in direct contrast to the modularity of mind approach, which claims that there are no external cognitive influences on perception (Fodor, 1983; Firestone & Scholl, 2016-a). For example, take the Ebbinghaus illusion (Figure 1). Most people have the percept that the orange center circle on the right is larger than the orange center circle on the left. However, they are actually the same size. With that said, knowing they are the same size does not make them actually appear the same size. This means that this visual

illusion is cognitively impenetrable. Some researchers question whether effects of action on perception challenge this notion of vision as being cognitively impenetrable.



*Figure 1.* The Ebbinghaus illusion. The center orange circles are the same size, however most people will have the percept that the center circle on the right is actually larger than the center circle on the left.

There have been a number of alternative suggestions for what may be driving action-specific effects. In cases for which estimation was made from memory (such as in the softball study), an alternative is that memory, rather than perception, is affected (Cooper, Sterling, Bacon, & Bridgeman, 2012). Said another way, perhaps it was not the case that softball players who were hitting better actually saw the ball as larger than did their less successful counterparts. It could have been, according to this alternative explanation, that they simply remembered it as being larger, again with no difference in actual perception. This is supported by, for example, literature which suggests memory can be easily manipulated to affect judgments of speed (Loftus & Palmer, 1974).

In other cases, perceivers were required to convert their perceptual experience into a report-able unit (e.g. feet). While verbal reports have been shown to be an effective method of ascertaining perceptions, they are also inherently susceptible to many kinds of biases

(Pagano & Isenhower, 2008; Poulton, 1979). Additionally, people do not have much experience in their day-to-day lives with converting their percept into units of measurement such as feet, yards, meters, etc. Some researchers suggest that these effects are nothing more than a response bias or judgment-related processes (Durgin, DeWald, Lechich, Li, & Ontiveros, 2011; Durgin et al., 2009; Firestone & Scholl, 2014, 2016-a). Perhaps those who hit the softball more successfully believed that they did so because the ball was bigger. Or maybe, participants were able to intuit the hypothesized pattern of results, and provided responses that were in line with experimenter expectations. The task of Action-Specific researchers is to establish whether these alternative explanations could be responsible for the reported effects. The importance of doing so is critical, not only because the participants' perceptions cannot be directly observed, but because this represents a challenge to the status quo of modularity. On one hand, these effects may simply be response bias, representing nothing more than an output level effect. On the other hand, these effects may truly be perceptual which would require reconsideration of the mechanisms which underlie perception.

Durgin, Klein, Spiegel, Stawser, & Williams (2012) have suggested that the social nature of experiments lead to participants adjusting their responses, and thus, the effects found are not perceptual, but are response bias. According to this account, in order to comply with the experimenter demands, the participant gives responses in line with what they think the experimenter wants the results to be. Therefore, it could be the case that instead of actually perceiving the distances up the hill to be farther, they are instead saying they appear farther to comply with the social demands. There have been a number of strategies to differentiate between perceptual effects and response bias. One example is to

directly measure whether the participants can even intuit the hypothesized results. Recently published work suggests that for the action-specific effect that distances up hills look farther than distances on the flat ground (cf. Stefanucci et al., 2005), only 25% of participants were able to intuit the hypothesized results, and over twice as many participants selected the *opposite* response (Tenhundfeld & Witt, 2017)! Another strategy is to use action-based measures or responses that do not involve conversion into units such as feet. For example, the effect that distances up a hill look farther has also been shown when performing a visual matching task for which participants adjusted a cone on the flat ground to be equidistance away from themselves as a cone up the hill, and also when perceived distance was measured using a blindwalking task (Tenhundfeld & Witt, 2017). Even when asked to blind walk the perceived distances, a measure lauded for its accuracy (Loomis, Da Silva, Fujita, & Fukusima, 1992), participants still reliably show these perceptual distortions (Tenhundfeld & Witt, 2017). Relying on measurements such as online comparisons of two stimuli, and action based measures, provides a compelling case that perhaps these effects are perceptual, rather than memory or percept-to-unit-conversion based.

For many action-specific effects, there is not yet sufficient research to determine whether the effect is due to a genuine difference in perception. However, with respect to one action-specific effect, there is sufficient evidence to rule out nearly all non-perceptual alternative explanations (Witt, 2017). To do so, Witt (2017) reviewed research using what will be henceforth referred to as the Pong task, for its similarities to the 1972 video game, Pong. In the task, participants are tasked with blocking a ball as it moves across the screen. The task of blocking the ball is made more difficult on some trials because the paddle is

small, and is made easier on other trials when the paddle is large. Participants are able to block approximately 45% of the balls with the small paddle, while they are able to over double that with the large paddle, blocking about 90%. The data has traditionally shown that when participants are given a small paddle with which to block the ball, the ball appears to move faster than it does when using a larger paddle. This discrepancy in perceived speed is theoretically driven by the difference in difficulty with each paddle. This difference between perceived speeds will be referred to as the Pong effect.

Given the robustness of the effect and level of experimental control over the Pong task, it provides an excellent platform with which to test influences on perception, including the challenges leveled by Firestone & Scholl (2016-a). As one example, even when feedback is given about the accuracy of participants' speed judgments, the Pong effect still emerged (King, Tenhundfeld, & Witt, 2017). As another example, after completing the Pong task, participants were asked to intuit the study hypothesis. Only 25% guessed correctly. And those who did not guess even with specific prompting showed the same magnitude of Pong effect as those who did guess correctly (Witt, Tenhundfeld, & Tymoski, in press). Given that the Pong effect still emerged even when controlling for other possible explanations, the evidence demonstrates that this particular effect likely demonstrates a true effect of action on perception (Witt, 2017; Witt, Sugovic, Tenhundfeld, & King, 2016).

However, this has not proven sufficient to quell the concerns of critics (Firestone & Scholl, 2016-b). Given the importance for theories of perception to differentiate perception from response bias and given that the empirical differentiation of the two is so challenging, it is imperative to continue to explore the two possible outcomes using a variety of strategies. The current experiments aimed to provide further insight into this important

but challenging problem by exploring perceptual differences when outcomes are maintained but action is removed from the task.

Participants will play the Pong task with the small and big paddles and, interspersed with these trials, there will also be trials for which the paddle controls itself (i.e. is automated). According to a response bias account of the Pong effect, participants infer that some aspect of the task is supposed to influence their speed judgments, and modify their judgments accordingly. For example, when the automated paddle is small, they might infer that is supposed to make the ball look faster and respond accordingly. Or if the ball is frequently blocked, they might infer that is supposed to make the ball look slower and respond accordingly.

According to a perceptual account, the size of the paddle is irrelevant except as it relates to performance (Witt & Sugovic, 2012). Therefore, if participants report the ball as faster when the paddle is small even if the paddle successfully blocks the ball on every trial, this would be strong evidence in favor of a response bias rather than a perceptual effect. On the other hand, if participants report the ball as moving slower when the automated paddle is small but very successful, this is consistent with both a response bias and a perceptual account.

A third possible outcome is that neither paddle size nor ball blocking performance impacts perception when the paddle is automated. This would be consistent with a perceptual explanation because in the case of an automated paddle, there is no action associated with the outcome. The paddle is moving itself, rather than the person acting with the paddle. If perception is truly action-specific, removing the action should eliminate the effect of the outcome on perception.

## Chapter 3: Experiment 1

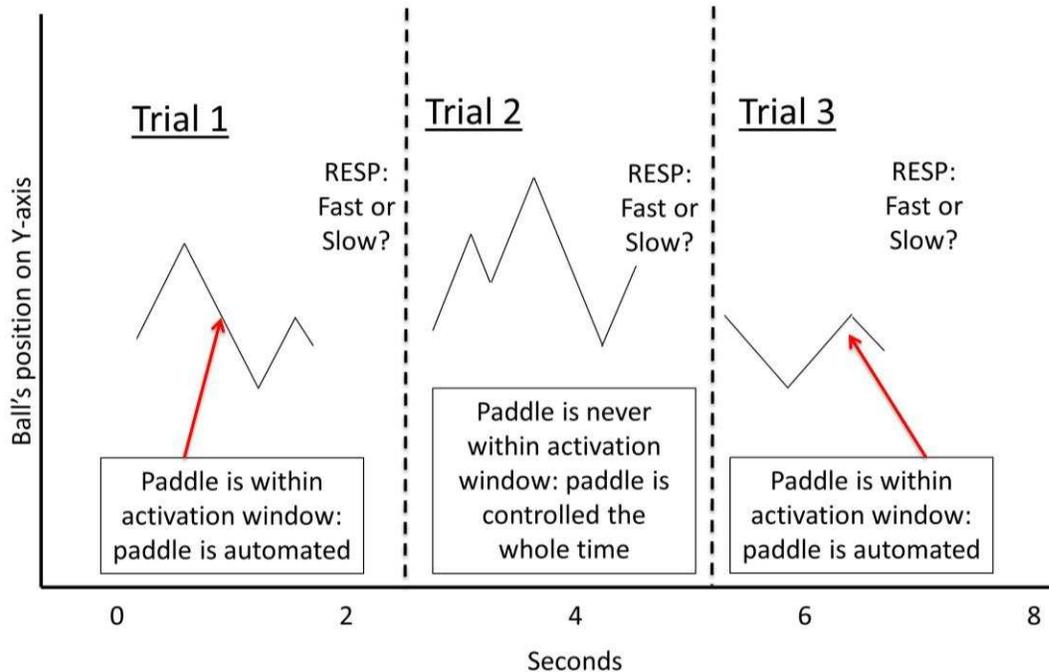
In this experiment I automated the blocking of the ball, in order to evaluate whether participants responses would be in line with a response bias based on the paddle size, or if it would be performance based, which could be accounted for with either a response bias based on performance, or the perceptual account. If participants respond according to a paddle size bias, the traditional Pong effect should still emerge even when there are no differences in performance. However, if participants respond according to a performance bias, there should be no difference in their PSEs when there is no difference in performance. However, as noted above, that result could also be indicative of a perceptual effect as well.

### Method

**Participants.** I recruited 27 undergraduates from the available research pool at Colorado State University.

**Stimulus and Apparatus.** The experiment was conducted on a 19" desktop display. The background was black. There was a 1.6 cm white circle, which served as the ball, and a 0.8 cm wide white bar that extended the length of the screen, vertically. Within that bar were two small horizontal black bars that were separated by either 1.9, 5.8, or 11.8 cm. The area in between the two horizontal black lines served as the paddle with which the participants tried to block the ball (Figure 3). The paddle's vertical placement was controlled by a joystick that participants moved. The ball moved at 1 of 6 speeds ranging from 26 to 67 cm/s from left to right (taking between 1724 and 493 ms, respectively, to reach the paddle) along a diagonal, changing directions randomly (for a visual

representation of trials, see Figure 2). The angle of the direction change remained constant within an individual trial, but varied by speed between trials (e.g. the angle was more obtuse for faster speeds and more acute for slower speeds).



*Figure 2.* A representation of trials in the automation condition. The ball's position is plotted over time. The red arrows indicate a point at which the automation might have engaged, but does not represent the point wherein it could or should have engaged. The length of each trial varied, but was in the same ball park as what is represented here. Everything was the same for the control condition except for the fact that at no point was automation engaged, regardless of paddle position.

**Procedure.** At the beginning of the experiment, participants were trained on the two anchor speeds, one slow (18 cm/s) and one fast (74 cm/s). The ball moved linearly across the screen at each of these anchor speeds three times (six in total), in a random order, without a paddle, and participants were told (via text on the screen) either “This is slow” or “This is fast”. Following this initial exposure, participants were exposed to the anchor speeds again, but this time the participant was not told which speed would be

shown and instead had to respond as to which speed was just demonstrated. Each anchor speed was shown three times (for a total of six additional trials), the order of which was randomized.



*Figure 3.* A representation of the display used for the experiment. The white circle served as the ball which bounced across the screen. The white area between the two horizontal black lines represented the paddle with which the participant had to try and block the ball. The paddle length varied between trials.

They were then asked to indicate if the ball moved at the slow or fast speed, and participants entered their responses by pressing the corresponding buttons on the joystick. They were given feedback following each response via text on the screen. At the end of the training, they were given the opportunity to restart the training if they wanted to, which no one did.

After the training of the anchor speeds, participants began one of two conditions (also known as the ‘start condition’), the order of which was counterbalanced across participants. In one condition, henceforth known as the no automation condition, participants completed 144 trials as has been done in previous studies. This served as a

replication of previous pong experiments. On each of these trials, participants were presented with one of the three paddle sizes, and the ball moved at one of the six speeds. The order and combination of these presentations was random, but each possible combination appeared an equal number of times. After each trial, participants responded as to whether the ball moved *more* like the slow speed or *more* like the fast anchor speed that they were trained on. They made their response by pressing the corresponding button on the joystick. In the other condition, henceforth known as the automation condition, participants were tasked with starting each trial by controlling the paddle as they did in the no automation condition. In this condition, however, the computer took over control of the paddle when a certain condition was met. The condition was that participants had to align the center of the paddle within 1.9 cm of the vertical placement of the ball, any time after the ball had traveled approximately  $\frac{1}{4}$  of the screen's width (taking between approximately 180 – 625 ms, depending on ball speed). For example, after the ball traveled at least  $\frac{1}{4}$  of its distance along the x-axis, if the center of the ball were at 20 cm from the bottom of the screen, the center of the paddle would have to be between 18.1 cm and 21.9 cm from the bottom in order for the automation to take over. Once the automation took over, the paddle moved up and down in perfect synchrony with the ball's position, and no longer changed position with movement of the joystick. This led the paddle to block the ball every time the automation was employed. Because of the design, if the ball was blocked, the automation was engaged. It was not possible to block the ball without the automation engaging. Following each trial, participants were asked to indicate if the ball moved more like the slow speed or more like the fast speed, regardless of whether automation was employed.

## Results

Speed judgments were submitted to a binary logistic regression in order to calculate the point of subjective equality (PSE) for each participant for each paddle length in each of the automation conditions. PSEs were then included in a boxplot analysis to examine for outliers, however none met my a priori classification of one score that was three times greater than the interquartile range (IQR) or two separate scores greater than one-and-a-half times the IQR. Thus, all participants were included in the analyses.

I ran a repeated measures ANOVA with paddle length and automation condition as within-subjects factors, and starting condition as between-subjects. There was a significant overall main effect for paddle length on PSE estimations (the Pong effect) such that the larger the paddle, the slower the ball appeared to move,  $F(2, 50) = 21.38, p < .001, \eta_p^2 = .46$ . The interaction between paddle length and automation was trending towards significance,  $F(2, 50) = 2.43, p = .099, \eta_p^2 = .09$ . Given the large variability I noticed in the boxplots, I replotted the data using difference scores (PSE with the big paddle – PSE with the small paddle) for each condition and found 2 participants with difference scores 1.5 times greater than the IQR. When these participants were excluded, the interaction between paddle length and automation was significant,  $F(2, 46) = 4.395, p = .018, \eta_p^2 = .160$ . The remaining analyses were similar regardless of these participants inclusion, so results are reported having kept their data in the analysis. There was no significant interaction between paddle size and the start condition,  $F(2, 50) = 0.129, p = .879, \eta_p^2 = .01$ . There was no significant main effect for the automation  $F(1, 25) = .744, p = .396, \eta_p^2 = .03$ , nor was there an interaction between automation and start condition,  $F(1, 25) = .023, p = .881, \eta_p^2 = .00$ . Difference scores were calculated between the big paddle PSEs and the

small paddle PSEs to quantify the magnitude of the Pong effect. Those difference scores were then compared between the automation condition and the no automation condition for both of the start conditions. There was a significant effect between the automation difference score and the no automation difference score when participants started with automation,  $t(13) = 2.318, p = .037, 95\% \text{ CI } [0.007, 0.208]$ . There was no significant effect between the automation difference score and the no automation difference score when participants started with no automation,  $t(12) = 0.726, p = .482, 95\% \text{ CI } [-0.089, 0.177]$ .

To further explore the data, I analyzed the pong effect in the automation conditions for each of the start conditions separately. The pong effect was significant when using automation for the no automation start condition,  $F(2, 26) = 4.19, p = .027, \eta_p^2 = .26$ , and significant when using automation for the automation start condition,  $F(2, 26) = 5.59, p = .010, \eta_p^2 = .30$ . In both cases, the big paddle resulted in a PSE that was 9.9% and 6.6% slower, respectively, than the small paddle (Figure 4).

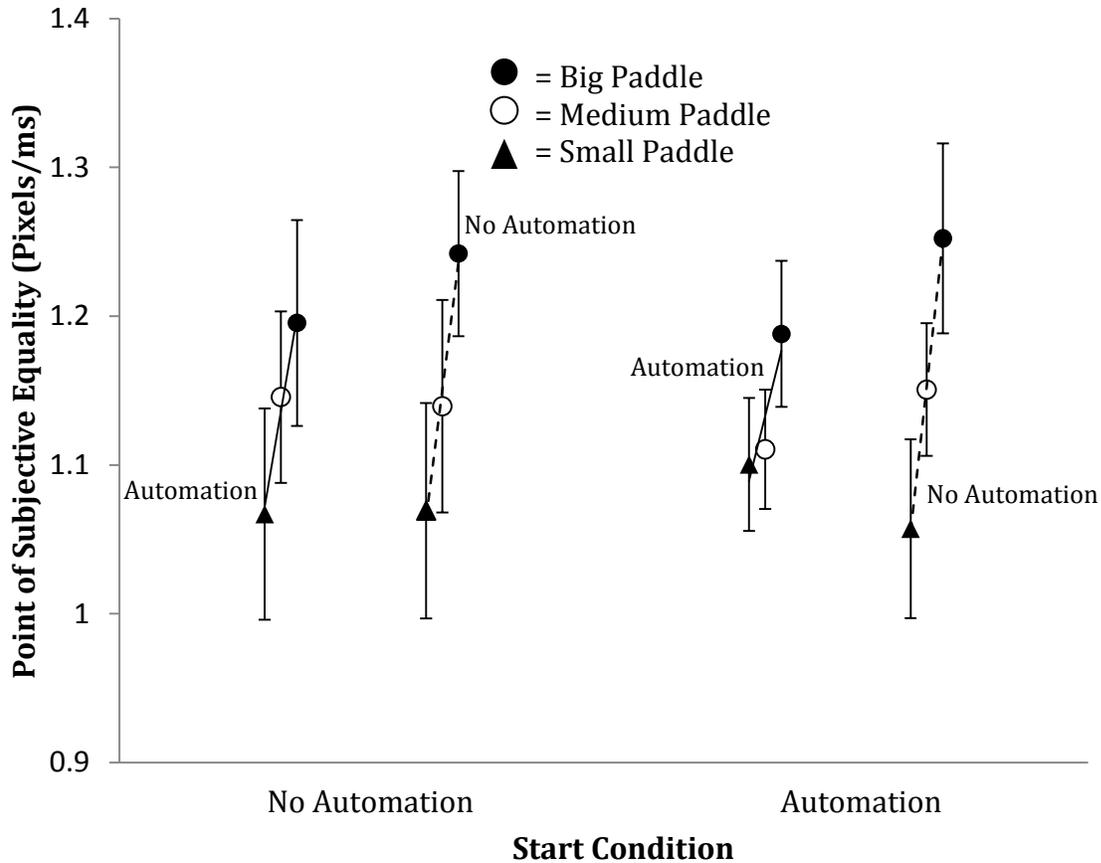


Figure 4. Point of subjective equality as broken down by paddle length, automation condition, and starting condition. Error bars are 95% confidence intervals, calculated within subjects.<sup>1</sup>

In order to verify that the automation had a significant effect on performance (as was theorized), I examined blocking performance between automation and no automation conditions. I ran a repeated measures ANOVA with automation condition and paddle length as within-subjects factors and starting condition as a between subjects factor. There was a significant main effect for the automation condition on ball blocking performance;  $F(1, 25) = 725.55, p < .001, \eta_p^2 = .97$  (Figure 5).

<sup>1</sup> Results are graphed in pixels per millisecond to aid in comparison within condition. For the values of each PSE in centimeters per second, to more easily compare this and all subsequent results to previous data, see Appendix A.

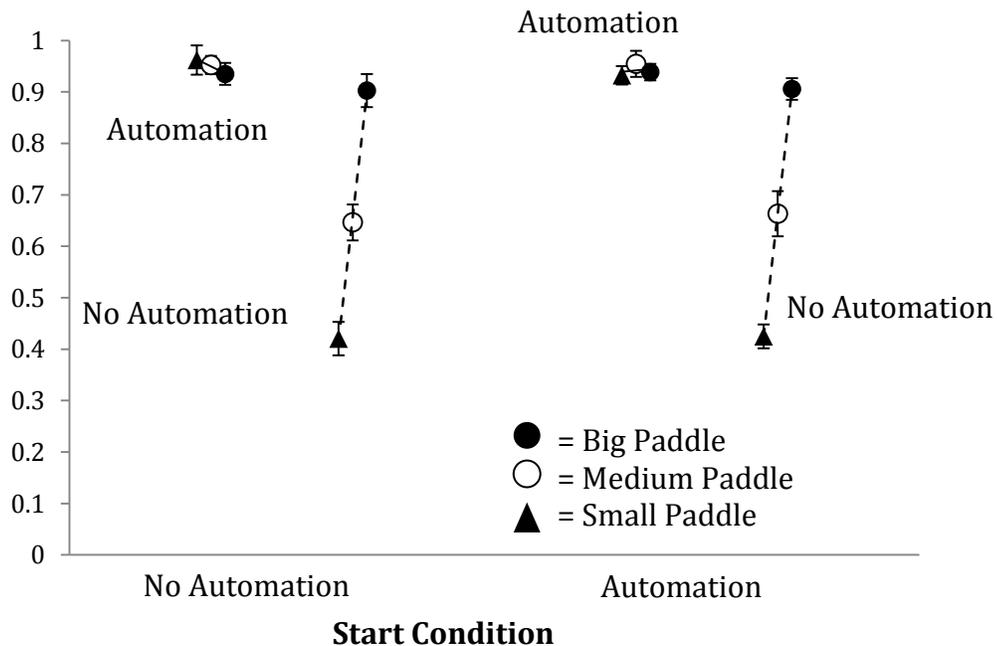


Figure 5. Proportion of balls blocked by paddle length, automation condition and starting condition. Error bars are 95% confidence intervals, calculated within subjects.

## Discussion

The fact that the Pong effect remained in the automation conditions, despite near perfect performance across all paddle sizes, is somewhat remarkable. Remember, the claim from action-specific researchers has been that wherein a change in performance is observed, there too should be a change in perception (e.g. Witt, 2011-a). Again, the conditions that engaged the automation were the exact same between paddle lengths, such that it was equally easy to engage with each paddle. Therefore, according to the action-specific account of perception, there should have been no difference in perception.

It is worth noting that the magnitude of the Pong effect in the automation condition, especially for those who started with automation, was smaller. However, the action-specific account would have predicted that there should be no difference in PSEs, not simply a

reduction. This result seems to suggest that a response bias explanation, based on the performance of each paddle, does not explain prior research, while a response bias explanation based on the size of each paddle might. However, it also suggests that perhaps the effect is not perceptual given that there was no difference in the difficulty between paddles, and yet a Pong effect still emerged.

With that said, and of particular note, is the fact that participants being required to act immediately preceding the automation taking over, may have contributed to these results. Considering the impact of intention and anticipation on the motor simulation of tasks (e.g. Witt & Proffitt, 2008), perhaps the fact that participants were starting the trial by acting led to an intention/motor simulation of completing each trial, irrespective of the automated boost in performance. I have chosen to use the definition of motor simulation as is discussed by Witt & Proffitt (2008): "...a motor simulation is the imagining of an action, either covertly or explicitly, without necessarily executing the action" and therefore "are future-oriented and have access to the outcome of anticipated action." If a motor simulation was engaged in the automation condition, then the fact that the end result (i.e. blocking) was automated may not have mattered in the formulation of their percept.

Additionally, the fact that there was a significant difference between the magnitude of the Pong effect for automation and no automation in the automation start condition, (such that the Pong effect in the automated trials when starting with the automation, was smaller than the effect during the no automation trials when starting with the automation condition), and no significant difference between the two magnitudes in the no automation start condition, may lend extra support to the motor simulation explanation for the results. These results might suggest that there was an improper calibration to the inherent

difficulty of the task (i.e. if it were performed without automation). For example, participants may have anticipated the small paddle would have been easier and the big paddle harder than they actually proved to be when the participant did not have automation. Thus, their motor simulation accounted for this anticipated difficulty in formulating their percept. When they were given actual information about the difficulty of the task (i.e. through performing the task without automation), their anticipated difficulty was accurately calibrated. While I can certainly not say definitively that this miscalibration was driving these results, the inference is somewhat bolstered by the fact that there was no difference between the Pong effect of the automation and no automation conditions when the participants started with no automation. In the no automation start condition, the participants started with the trials that would help calibrate them to the actual difficulty of the task. This calibration to the actual difficulty should have then theoretically led to a similar Pong effect in the automation condition if calibration to the difficulty were at play in the motor simulation explanation. That is what was observed. With that said, because participants were not asked to indicate the expected difficulty of the task before hand, this is purely speculative. Finally, despite this reduction in size of the effect for the automation trials in which they started with automation, the pong effect remained significant.

From these results, it would be reasonable to say that, at least in this specific condition, that the Pong effect is likely not driven by a response bias based on the performance. If it had been, perceived speed would have been the same across the paddle sizes for the automation condition, given the similar performance. What is not clear, however, is if that means the Pong effect may be driven by a response bias based on paddle length.

In summary, these results indicate that the Pong effect is likely not due to a response bias based on performance given that equal levels of performance did not produce perceptions of equal speeds. However, as discussed above, it is not clear if the emergent Pong effect was due to a motor simulation, or because of a response bias based on the paddle sizes. Said another way, the fact that even on trials in which the participants had automation they were required to act for at least the first quarter of the trial, could have elicited a perceptual distortion, as they were disregarding the actual performance and focusing on their inherent ability to act. Alternatively, the observed Pong effect in the automation trials could have been driven by a response bias caused by the differences in paddle size. In order to address these concerns, and to more directly test whether the Pong effect is driven by response bias or perception, I ran Experiment 2.

## Chapter 4: Experiment 2

This experiment directly manipulated performance without the confound of control over the automated paddle. This allowed me to examine whether a small automated paddle that had the performance of the big controlled paddle (i.e. no automation), would show a pattern of results in line with the small paddle (which is the same size) or the large paddle (which has the same performance). If the PSEs when using the small automated paddle are similar to the PSEs when playing with the small controlled paddle (i.e. no automation), this would suggest that participants rely on the size of the paddle to make their speed estimations. If the PSEs when using the small automated paddle are similar to the PSEs when playing with the big controlled paddle, this would suggest that participants might produce a response bias based on the performance of the paddle. However, if the action-specific account of perception is right, by removing the ability of the participant to act, there should also be no reliance on either their own inherent ability, nor the performance of the automated paddle. Said another way, if action affects perception, when there is no action, there should be no effects of action on perception.

### Method

**Participants.** I recruited 30 undergraduates from the available research pool at Colorado State University.

**Stimuli & Apparatus.** The stimuli and apparatus are the same as in Experiment 1, with the following exceptions. There were only two paddle sizes: small and big. The paddle used to block the ball was either red or blue. This coloring allowed for distinctions to be made between trials in which the participants had control, versus trials in which the paddle

was automated. The color for each condition remained constant for the duration of the trials for each participant, but was randomly assigned between participants. Said another way, the first participant would always see a red paddle on those trials in which s/he had control over the paddle, and a blue paddle anytime the paddle was automated. The second participant would always see a blue paddle on those trials in which s/he had control over the paddle, and a red paddle any time the paddle was automated.

Additionally, in order to have greater sensitivity at the most ambiguous speeds, half of the trials at the slowest and fastest speeds were changed to be at the two middle speeds (of the six total). This did not change the number of trials, nor the range of experimental speeds, but simply made the theoretically most ambiguous (in relation to their classification of more like the slow or more like the fast anchors) speeds appear more frequently. This increased my ability to detect differences at these critical mid-range points.

Because of a subjective experience that the paddle could still be controlled even when it was automated, I added a survey at the end of the experiment in order to ascertain whether the subjects believed they could control it. The survey said: "After a few trials, did you feel that you had any control over the [automated paddle color] automated paddle? Enter the number that best corresponds to your experience:

1. Yes I felt I could control it throughout the experiment
2. Sometimes I felt I could have some control over it
3. No, I almost never felt that I could control it"

**Procedure.** The procedure is the same as in Experiment 1 with the following exceptions. In this experiment the control versus automation condition remained constant

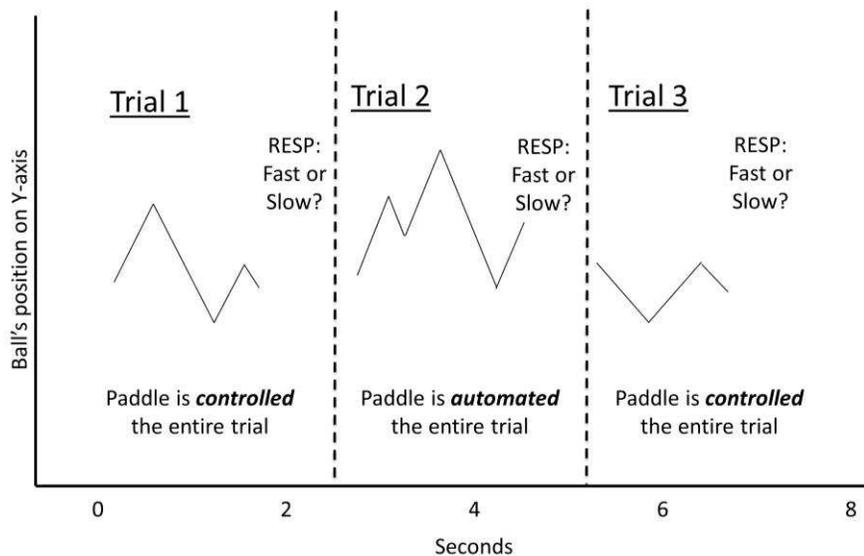
for the duration of the trial. Whereas in Experiment 1 participants were required to act in a certain way that would engage the automation, in Experiment 2, trials were randomly either fully automated or fully controlled. There were 288 experimental trials broken into three types of trials.

The first trial type was the typical small paddle trial wherein the ball bounced across the screen and participants were tasked with trying to block the ball by moving a joystick to move the corresponding virtual paddle. The second trial type was the same as the first, except this time the paddle was 5 times longer (again participants were in control). These two trial types will be referred to the small controlled paddle and the big controlled paddle respectively. For the third trial type, the paddle was the same size as the small controlled paddle, however this paddle was automated, which made it as good at blocking the ball as an average participant with the big controlled paddle (~95% of balls successfully blocked). This is the examination into performance; by making the small automated paddle functionally as efficient at blocking as is the large controlled paddle, I can examine the degree to which the response of participants is more in line with the size of the paddle, or with the performance of the paddle. For trials of the third type, the automation was always employed and thus there was no possibility for participants to control the paddle.

The movement of the automated paddle was non-biological; whereas there is a continuous movement when the participant controls the paddle, with automation the paddle jumps every so often (the frequency was determined a priori in order to match mean performance of the large controlled paddle). The paddle still tracked the movement of the ball but in a less fluid, more disjointed manner. This was done in order to isolate the

performance aspect. Whereas automation that behaves more like a human may be more likely to be embodied, automation that does not act like a human, nor provide the ability to act, should not. After each trial, participants were asked to indicate if the ball moved more like the fast anchor speed, or more like the slow anchor speed. The researcher would watch for the first several automated trials to ensure that the participant was not trying to control the paddle. If the participant was trying to control the automated paddle (by moving the joystick as if they were in control), the researcher would kindly remind them that when the paddle whichever color designated automation, the participant did not have to control it (for a visual representation of example trials, see Figure 6).

After all 288 trials, the participants were shown the survey and asked to respond accordingly.



*Figure 6.* Visual representation of example trials. The ball's position is plotted over time. Whether the paddle was automated or controlled was indicated by paddle color, and remained the same over the duration of the trial, but would vary between trials. The length of each trial varied, but was in the same ball park as what is represented here.

## Results

Speed judgments were submitted to a binary logistic regression in order to calculate the point of subjective equality (PSE) for each participant for each of the three paddle conditions. In order to avoid any confounds that a belief of control may have had, any participant who responded that they felt they could control the automated paddle throughout the experiment were excluded. This led to the exclusion of three participants. The remaining 27 PSEs were then submitted to a boxplot analysis to look for outliers. An a priori criterion was set such that any participant with one score that was three times greater than the interquartile range (IQR) or two separate scores greater than one-and-a-half times the IQR. Given this cut off, one additional participant was excluded. Finally, one participant had blocking performance that suggests they were not paying attention (only 25% blocked with the big paddle). For that reason, they were also excluded from further analyses, leaving 25 total participants.

The automated paddle was programmed to successfully block the ball approximately 95% of the time. Results showed that performance with the automated paddle slightly exceeded that target,  $M = 96.9\%$ ,  $SD = 1.70\%$ . Performance was significantly better with the small automated paddle than the big controlled paddle ( $M = 89\%$ ,  $SD = 5.8\%$ ),  $t(24) = 6.95$ ,  $p < .001$ , 95% CI [5.5, 10.2], and the small controlled paddle ( $M = 44.7\%$ ,  $SD = 7.7\%$ ),  $t(24) = 36.48$ ,  $p < .001$ , 95% CI [49.3, 55.2].

The purpose of this study was to see if the automated paddle induced a percept that was more in line with the small controlled paddle or the big controlled paddle. As shown in Figure 7, the automated paddle induced a PSE somewhere between the two control paddles.

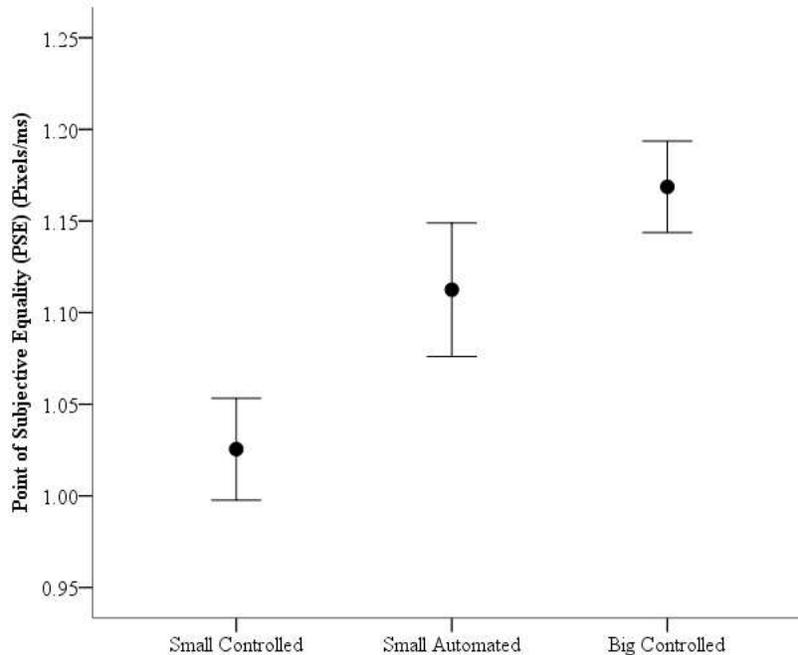


Figure 7. PSEs for the three paddle conditions. Error bars represent the 95% confidence interval calculated within-subjects.

A paired t-test between the small controlled paddle PSEs and the small automated paddle PSEs showed a significant difference,  $t(24) = 3.001$ ,  $p = .006$ , 95% CI [.027, .147]. There was also a significant difference between the small automated paddle and the big controlled paddle,  $t(24) = 2.073$ ,  $p = .049$ , 95% CI [.000, .112].

The difference between the PSE for the small automated paddle and the point of objective equality (POE) calculated as the average of the 6 test speeds (1.125 pixels/ms) was not significant,  $t(24) = 0.144$ ,  $p = .887$ , 95% CI [-.053, .0614]<sup>2</sup>.

<sup>2</sup> It is also possible to calculate the POE as a function of the average between the two initial anchor speeds, which would yield an average of 1.25 pixels/ms. However I believed that it was more appropriate to calculate the objective equality on the speeds they were classifying, rather than the classifiers (the anchor speeds) themselves.

## Discussion

The results of Experiment 1 indicated that even when there was no difference in performance, there was a difference in PSEs. This suggested that perhaps results could be explained with a response bias account. However, on each of those automated trials participants were required to control the paddle until it met certain preconditions for automation to take over. This may have inadvertently imposed a confound in the data such that participants were subsequently seeing the trial as if they would have continuous control over that paddle. Because of that, I ran the current experiment in which a small automated paddle was automated for the entirety of the trial. The automated paddle was the same size as the small controlled paddle, but had performance that actually proved to be better than the big controlled paddle.

If participants were responding in accordance with the *size* of the paddle, there should have been no difference between the PSEs when using the small controlled and the small automated paddles. If the participants were responding in accordance with the *performance*, there should have been no difference between the PSEs when using the big controlled and the small automated paddles. However, the action-specific account of perception would predict that because there was no action with the small automated paddle, neither the performance of the paddle, nor the participants' inherent abilities, should matter. Therefore the action-specific account would predict that the PSEs would be in line with the optically specified information represented by the POE.

Results indicated that the speed perceptions when using the small automated paddle were significantly different than when using the small controlled paddle. However, despite performing better on average with the small automated paddle than they did with

the big controlled paddle, the PSE for the big controlled paddle was significantly higher than the PSE for the small automated paddle. These results suggest that the Pong effect is likely not driven by a response bias in accordance with either the size or performance of the small automated paddle. Remember, that the demands which would have led to the response bias were still found in the present experiment. Said another way, if the Pong effect could be explained by a response bias, the present experiment should have led to the response bias still emerging. Both the paddle size and performance information were still there, and they were the same as in previous experiments. However, not only did a response bias not emerge, but the results are entirely consistent with the action-specific account of perception. Because no action was required, nor even possible, there should have been a complete discarding of action information, and thus a reliance on the optically specified information, which is what was found. The fact that the PSEs when using the small automated paddle were no different than the POE supports the claim that participants were relying on only the optically specified information to formulate their percept, and their subsequent responses.

With that said, and given concerns in the field surrounding replicability (Open Science Collaboration, 2012), I wanted to do a direct replication before moving on.

## Chapter 5: Experiment 3

In order to verify the pattern of results in Experiment 2, I ran a direct replication.

### Method

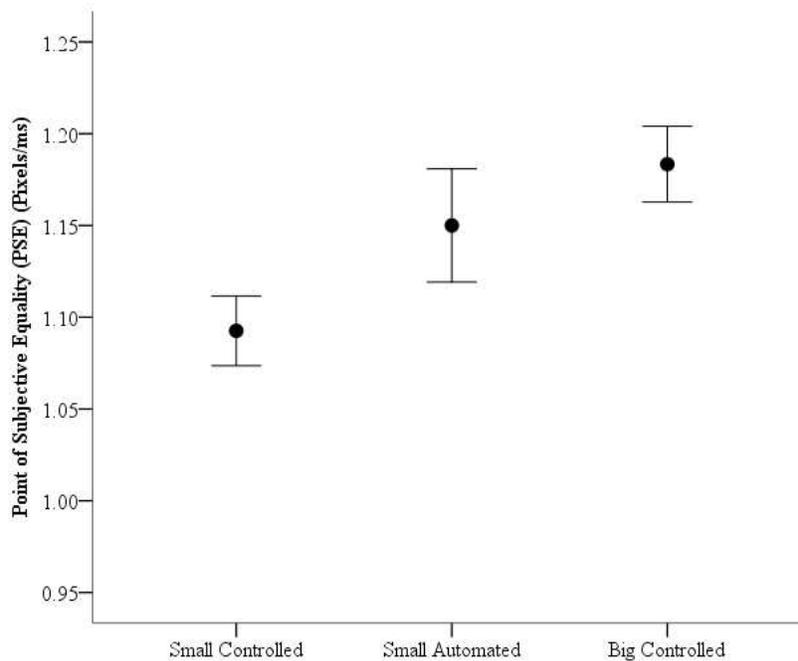
I recruited 27 undergraduates from the available research pool at Colorado State University. The stimuli, apparatus, and procedure were the same as in Experiment 2.

### Results

Speed judgments were submitted to a binary logistic regression in order to calculate the point of subjective equality (PSE) for each participant for each of the three paddle conditions. In order to avoid any confounds that a belief of control may have had, any participant who responded that they felt they could control the automated paddle throughout the experiment were excluded. This led to the exclusion of one participant. The remaining 26 PSEs were then submitted to a boxplot analysis to look for outliers. An a priori criterion was set such that any participant with one score that was three times greater than the interquartile range (IQR) or two separate scores greater than one-and-a-half times the IQR. Given this cut off, no additional participants were excluded.

There was a significant difference between PSEs for the small controlled paddle and the small automated paddle,  $t(25) = 2.522$ ,  $p = .018$ , 95% CI [.011, .104], replicating the finding in Experiment 2. The difference between the small automated paddle and the big controlled paddle was not significant this time, however,  $t(25) = 1.404$ ,  $p = .173$ , 95% CI [- .016, .082] (Figure 8). However, the trend was in the same direction. Given the small nature of this effect size, I combined the data between Experiments 2 and 3, which yielded

a highly significant difference between the small automated paddle and the big controlled paddle,  $t(50) = 8.098, p < .001, 95\% \text{ CI } [.063, .104]$ .



*Figure 8.* PSEs for the different paddle conditions. Error bars represent 95% confidence intervals calculated within-subjects.

The small automated PSE was not significantly different from the POE,  $t(25) = 1.197, p = .243, 95\% \text{ CI } [-.025, .094]$ .

## Discussion

These results serve to corroborate the findings of Experiment 2. However, the interpretation of the results is slightly ambiguous. On one hand, PSEs were slower with the automated small paddle than with the controlled small paddle, which further supports the claim that PSEs were not subject to a response bias based on paddle size. On the other hand, the PSEs when using the small automated paddle were not significantly different from the PSEs when using the big controlled paddle. However, the data still appear to be

trending towards a significant difference between the PSEs of the small automated paddle and those of the big controlled paddle. This is corroborated by the highly significant difference found when combining the data of Experiments 2 and 3. For that reason, results will be discussed in regards to a replication of the results of Experiment 2.

Again, that the PSEs were not in line with a response bias based on either the paddle size or paddle performance seems to suggest that the Pong effect is not driven by a response bias. The small automated paddle had the size of the small controlled paddle and the performance of the big controlled paddle. If a response bias could be driving the PSEs observed when using the small and big controlled paddles, it too should have driven the small automated paddle in one direction. The difference between the automated and controlled paddles is of action capability. Because there was no ability to act, the action-specific account of perception would suggest that there should be no incorporation of action information, thus relying on optically specified information to form the percept.

Consistent with the results of Experiment 2, there was no significant difference between the PSEs when using the small automated paddle, and the POE. Once again, this would seem to indicate that when the ability to act has been removed from the participant, they are discarding the performance and action information, thus relying on the optically specified information to formulate their percept.

However, one more study is needed in order to be able to determine if the Pong effect can be explained away with response biases, or if it represents a genuine effect of action on perception. Perhaps the results of Experiments 2 and 3 were due to an averaging of some participants' responses being representative of a response bias based on paddle size, with those whose responses were representative of a response bias based on

performance. Additionally, perhaps the fact that these two sources of information that a participant could be using were contradictory, could have led to them generating a response that was in the middle (which is what was observed).

To evaluate this potential explanation, I changed the performance of the small automated paddle to be in line with the small controlled paddle. If it turns out to be the case that the cause of the small automated PSEs in Experiments 2 and 3 is an averaging of the two response biases, by setting the small automated paddle's performance to be where the small controlled paddles performance is, there should be a shift in the PSE of the small automated paddle to be in line with the small controlled paddle. Said another way, if the PSEs of the small automated paddle in Experiments 2 and 3 were because of the two response bias accounts contradicting each other, by setting them to predict the same results, the PSEs of the small automated paddle should be in line with the response bias prediction. However, if the results are caused by a disregard for action information the pattern of results should remain the same. Thus, I ran Experiment 4.

## Chapter 6: Experiment 4

Because the pattern of results in Experiments 2 and 3 could have still theoretically been explained by a response bias account, I decided to make it such that the two response biases would predict the same result. By changing the performance of the small automated paddle to be the same as the small controlled paddle, whether participants were relying on the paddle size or paddle performance to generate a response, the results should be the same. However, if the action-specific account is to be believed, there should be a discarding of action information all together when the ability to act has been removed. Therefore this experiment is able to pit the claims of response bias against the theory of action-specific perception. With the response bias account predicting the PSEs of the small automated and small controlled paddles being no different, and the action-specific account predicting they will be, this experiment will serve to concretely pit the two accounts against each other like has never been done before.

### Method

**Participants.** I recruited 30 undergraduates from the available research pool at Colorado State University.

**Stimuli & Apparatus.** The stimuli and apparatus was the same as in Experiment 3.

**Procedure.** The procedure was the exact same as in Experiment 3 with the following exceptions. In this experiment the blocking performance with the small automated paddle was set to be approximately 45%, to mimic the performance with the small paddle. To achieve this, the automated paddle was programmed to update its position every 188 ms with a random amount of y-axis positioning noise. This update rate

and noise in the positioning of the paddle led to the paddle not always being in the necessary position to block the ball when the ball would reach the position of the paddle along the x-axis, thus leading to a blocking rate of 45%.

## Results

As a manipulation check, I examined ball blocking performance for the three conditions. Performance with the small automated paddle ( $M = 44.0\%$ ,  $SD = 4.9\%$ ) was not significantly different from the small controlled paddle ( $M = 42.9\%$ ,  $SD = 8.8\%$ ),  $t(27) = 1.045$ ,  $p = .305$ , 95% CI [-2.0, 6.2]. Performance with the small automated paddle was significantly worse than performance with the big controlled paddle ( $M = 89.3\%$ ,  $SD = 5.3\%$ ),  $t(27) = 30.070$ ,  $p < .001$ , 95% CI [42.2, 48.3]. The small controlled paddle performance was also significantly worse than the big controlled paddle,  $t(27) = 35.233$ ,  $p < .001$ , 95% CI [44.6, 50.1].

Speed judgments were submitted to a binary logistic regression in order to calculate the point of subjective equality (PSE) for each participant for each of the three paddle conditions. In order to avoid any confounds that a belief of control may have had, any participant who responded that they felt they could control the automated paddle throughout the experiment were excluded. This led to the exclusion of two participants. The remaining 28 PSEs were then submitted to a boxplot analysis to look for outliers. An a priori criterion was set such that any participant with one score that was three times greater than the interquartile range (IQR) or two separate scores greater than one-and-a-half times the IQR. Given this cut off, no additional participants were excluded.

Pairwise comparisons indicated a significant difference between the small controlled paddle and the small automated paddle  $t(27) = 2.800$ ,  $p = .009$ , 95% CI [.056,

.363], and a significant difference between the small automated paddle and the big controlled paddle,  $t(27) = 2.894, p = .007, 95\% \text{ CI } [.020, .119]$  (Figure 9).

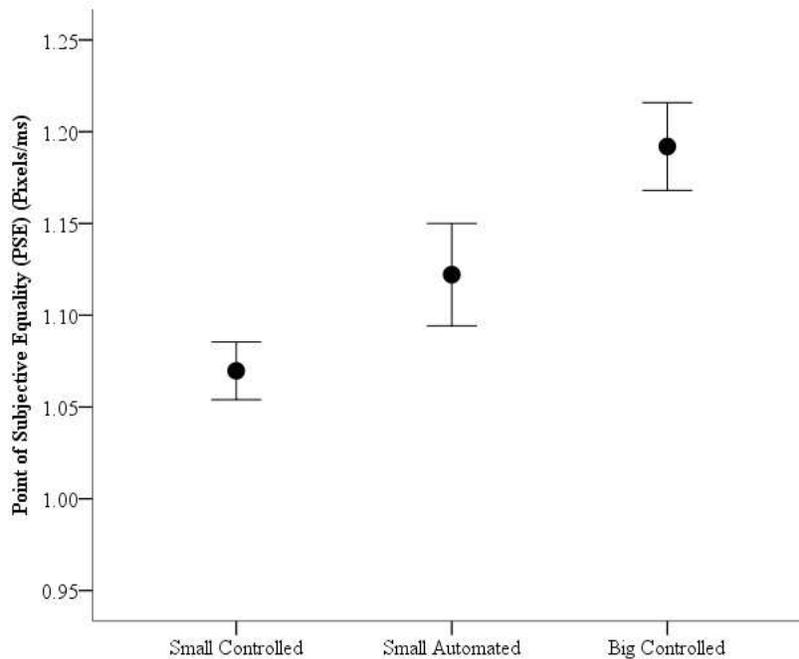


Figure 9. PSEs for the different paddle conditions. Error bars represent 95% confidence intervals, calculated within-subjects.

Additionally, there was no significant difference between the PSE for the small automated paddle in Experiment 4 versus in Experiments 2/3 (I collapsed between experiments for this analysis),  $t(77) = 1.112, p = .269, 95\% \text{ CI } [-.028, .098]$ . There was no significant difference between the PSE for the small automated paddle in Experiment 4 and the POE,  $t(27) = .706, p = .486, 95\% \text{ CI } [-.030, .060]$ .

## Discussion

In this experiment I manipulated performance such that the small automated paddle would have the same performance as the small controlled paddle. If an averaging of the potential response biases drove the results of Experiments 2 and 3, in this experiment the

PSE with the small automated paddle should have been similar to the PSE with the small controlled paddle. Said another way, because in this experiment the small automated paddle was no different in performance or size when compared to the small controlled paddle, regardless of the source for a response bias the PSEs for the small automated and small controlled paddles should have been the same. However, if a discarding of action and performance information all together drove the results of Experiments 2 and 3, the results of this experiment should have been no different than those of Experiments 2 and 3.

Results indicate that not only was there a significant difference between the PSEs for the small controlled paddle and the small automated paddle, but that there was no significant difference in the PSEs when using the small automated paddle between Experiments 2 - 4. The fact that there was no significant difference between the PSEs when using the small automated paddle in Experiments 2 - 4 is additionally compelling. Regardless of the source of information driving a response bias (i.e. paddle size or performance), the results of this experiment (Experiment 4) should have been different if a response bias were able to account for the Pong effect.

These results, in tandem with the fact that the PSEs when using the small automated paddle in Experiments 2 - 4 were not significantly different than the POE, suggests the results in Experiments 2 - 4 were driven by a complete disregard for paddle performance information. Said another way, it seems as though instead of utilizing either the paddle's performance, or the paddle's size to formulate a response, which would have suggested the Pong effect was a response bias, the perceptual system was discarding action information, and instead just relying on the optically specified information. Additionally, this suggests

that a motor simulation of the task, rather than a response bias of paddle size in fact drove the results of Experiment 1 (see further discussion of this point in the General Discussion).

Implications for these results, as well as limitations and additional considerations will be discussed at length in the General Discussion. However, to preview the discussion, it appears as though the Pong effect cannot be explained with a response bias account. This, in addition to the abundance of other research controlling for alternative explanations seems to suggest that the Pong effect is an actual demonstration of action affecting perception. This is critical in that it demonstrates not that every reported action-specific effect is a genuine effect on perception, but rather that there *can* be effects of action on perception. By demonstrating a single example of where action affects perception, this requires that theories of vision incorporate consideration that the perceptual system is not modular (Fodor, 1983), but instead incorporates information about the body's ability to act.

## Chapter 7: General Discussion

### Perception versus Response Bias

Previous research has shown that our ability to act affects how we quite literally see our surrounding environment, a theory known as the action-specific perception (Proffitt, 2006; Witt, 2011-a). For example, research has found that individuals who do parkour see wall heights as being less tall than do their novice counterparts (Taylor et al., 2011). The suggestion from action-specific researchers is that because the skilled free runners (those who do parkour) are better able to scale walls, their perceptual system has taken into account this ability and distorts their perceptions accordingly. Whereas those who are novices at climbing up walls do not have the proper technique to do so and thus would be less successful in their attempts to scale walls of the same height, a reality that their perceptual system accounts for.

Action-specific results have been met with considerable controversy however. To many who ascribe to the idea that the perceptual system is modular, which is to say it is cognitively impenetrable (Firestone & Scholl, 2014; Firestone & Scholl, 2016-a), the idea that action information can affect perception is incongruent with the way the perceptual system works. According to the idea that perception is cognitively impenetrable, there should be no top-down influences on perception. That is to say that the perceptual system is isolated from all effects of cognition. The question as to whether these action-specific effects are in fact perceptual is paramount to our understanding of how the perceptual system works. Firestone and Scholl (2016-a) put it best: “The question of whether there are top-down effects of cognition on visual perception is one of the most foundational

questions that can be asked about what perception is and how it works...” As they go on to point out, this debate has ramifications beyond psychology, and into fields of philosophy, neuroscience, psychiatry, and aesthetics (Firestone & Scholl, 2016-a).

Because action-specific effects represent such a revolutionary view on “what perception is and how it works” (Firestone & Scholl, 2016-a), many have posited alternative explanations for effects that action-specific researchers call perceptual. One potential alternative explanation is that the Pong effect, and other action-specific effects like it, is simply a response bias (e.g. Durgin et al., 2009; Firestone & Scholl, 2016-a). Considerable work has been done to demonstrate that, at least in the Pong task, it seems as though the effects are in fact perceptual (i.e. Witt, 2017; Witt et al., 2016). For example, Witt et al. (2016) demonstrated that the Pong task likely represents a genuine effect of action on perception by examining how past research has controlled for alternative explanations. However, it is still theoretically possible that response biases exist given the nature of the design.

The response bias argument would go along the lines of: the manipulation of paddle size is obvious to participants and therefore, in an effort to be a good participant, they are responding based on what they think the hypothesized direction of results is. In this case there could be two potential response bias explanations for the results. First, because the paddle was small, it must be harder to use and therefore the participant would indicate that the ball was moving faster, even though it did not actually look like it was moving faster. The second possible response bias explanation would suggest that because the small paddle performs significantly worse than the big paddle, participants respond as if it is

moving faster even though, again, it does not actually look like it is moving faster. Both of these explanations predict the same pattern of results as the action-specific account.

Previous research has attempted to determine whether these effects are response bias or perceptual. King and colleagues (2017) provided feedback to participants. Their supposition was that by imposing a response bias (i.e. to be as accurate as possible), this should negate the effects of any other response bias. The idea is that if participants are adjusting their responses in order to be good participants, they will respond as accurately as possible given that it is the experimenter demand placed on them. The authors argued that if the Pong effect still emerged, it would be representative of a truly perceptual effect that was not caused by a response bias (King et al., 2017). If a response bias that imposed the demands of accuracy, could not be met, it would suggest that participants' were not able to overcome the compelling perceptual distortion. The Pong effect still emerged. However, it could be argued that perhaps the imposed response bias was not effective enough to overcome the inherent response biases of paddle size or performance. Therefore, the question remains, is the Pong effect a genuine demonstration of action's effect on perception, thereby helping reshape the entire way vision researchers conceptualize the perceptual system? Or, is the Pong effect nothing but a response bias dressed in the clothing of a perceptual effect by those who wish it so?

The present experiments address these response bias claims in an entirely unique way: by isolating and manipulating the different potential response bias explanations. If participants intuited a hypothesis based on the size of the paddle, results of Experiments 2 – 4 should have indicated no difference in PSEs between the small controlled and small automated paddles. Because these paddles were the same size, the same demands would

have existed to respond accordingly. If however the participants intuited a hypothesis based on performance and responded accordingly, results would have indicated no significant difference between the small automated paddle and the big controlled paddle in Experiments 2 and 3, and no difference between the small automated paddle and the small controlled paddle in Experiment 4. This is because, again, the performance was the same between those conditions. If a participant is biased in responding based on performance for either of the controlled paddles, they should have also been biased for responding based on performance for the automated paddle.

However, the results are not consistent with a response bias based on either paddle size, or performance. Whereas both response biases would predict that the PSEs when using the small automated paddle would be in line with either the small controlled or the big controlled paddles, results indicated that the PSEs were in between those of the small and big controlled paddles. One might claim that perhaps the reason the PSE was in the middle of the small and big controlled paddles in Experiments 2 – 4 is because there was a split with some participants intuiting a hypothesis based on paddle size with the other half intuiting a hypothesis based on performance. However, as you will notice, both of those hypotheses would have predicted that in Experiment 4 the PSEs when using the small automated paddle should have been no different than the PSEs when using the small controlled paddle. Importantly, in Experiment 4 the response bias account predicted that there should be no difference between the small controlled and the small automated paddle, whereas the action-specific account predicted that there should remain the difference that existed in Experiments 2 and 3. This difference between the PSEs of the small controlled and small automated paddle remained in Experiment 4.

Additionally compelling, the PSEs with the small automated paddle in Experiment 4 were not significantly different than those with the small automated paddle in Experiments 2 and 3, and none of them were significantly different from the point of objective equality (POE). This is exactly what the action-specific account of perception would predict. The action-specific account suggests that when the ability to act is removed, there would be no effect of action on perception. Therefore, it would not matter what the performance of the automated paddle is in Experiments 2 – 4 because it has no bearing on the individual. Thus, the perceptual system would discount this performance information, as well as any inherent action information, and rely on optically specified information to form the percept.

To reiterate, response biases that exist based on paddle size or paddle performance cannot explain the Pong effect. The demands placed on the participant to respond in accordance with one or the other when using the controlled paddles, were there when watching the automated paddle. The difference was in the participants' action capabilities. Whereas they could act with the controlled paddles, they could not act with the automated paddles. When the participants could act, they seemed to incorporate action information into their percept of ball speed. When the participants could not act, they discarded information about action and performance, relying instead on the optically specified information available to them. Not only have I found effects where effects should exist, but I have also critically found no effect where an effect should not exist. This represents another example of how Firestone and Scholl (2016-a) suggest researchers evaluate their perceptual claims of action's effect on perception.

My data also speak to another potential explanation for the Pong effect which would still be perceptual, but would not rely on action information to explain. This alternative has to do with the Leibowitz Hypothesis (Leibowitz, 1985). The Leibowitz Hypothesis says that larger objects appear to move slower than do smaller ones. For that reason, it is possible that the effect of balls appearing to move slower when using a large paddle, in lieu of a small paddle, could be attributed to some sort of an assimilation of ball speed to the perceived paddle speed. However, if that were the case, one would expect that in Experiments 2 – 4 the perceived ball speed when using the small paddle would have been the same regardless of whether that paddle was automated or not.

Additionally, it could have been the case that the reason the reason the PSEs with the small automated paddle were in between those when using the small controlled and the big controlled paddles was because of an averaging of the competing perceptual information. An assimilation to the perceived speed of the paddle would have resulted in the balls appearing to move faster when using the small paddle. However the performance information when using the small automated paddle would have resulted in the balls appearing to move slower. An averaging of these two perceptual effects could have explained the resulting PSEs. However, the fact that there was no difference in the pattern of results between Experiments 2 – 3 and those of Experiment 4 indicates that it is also likely not the case that the results were an averaging of perceptual information from the Leibowitz Hypothesis and action information.

While this data does not explicitly demonstrate that because the pattern of results was not consistent with a response bias or alternative perceptual explanation, that it then *must* be a genuine effect of action on perception, it does add a voice to the choir suggesting

that non-perceptual accounts of at least the Pong effect may be misguided. This is critical because if one is to believe that action genuinely can affect perception, this would then mean that the way in which researchers have thought about the perceptual system necessitates reconsideration. Perhaps the perceptual system is not cognitively impenetrable as is suggested by some of the most vehement critics (Firestone & Scholl, 2014; Firestone & Scholl, 2016-a). This would mean that we could no longer consider visual perception to be ‘modular’, but rather to be interconnected (Fodor, 1983). As mentioned above, this revolutionizes how psychologists, philosophers, neuroscientists, and others determine what not only what perception is and how it works, but what this means for long standing views about how the mind works.

The breadth of what top-down effects on perception exist still warrant the same level of scrutiny. Simply because I have shown that the Pong effect is not a response bias does not mean that all reported action-specific effects (or even all top-down effects on cognition) are not response biases. However, these results add that little extra bit of foundational stability to the claim that there *can* be top-down effects of action on perception. If we, as a field, can agree that one instance exists, this forces us to consider that others might as well, and thus will lead to a much broader understanding about the workings of the perceptual system!

### **A Role for Motor Simulation**

Given the past literature with tools, it seems as though one might have predicted that the automation’s performance could have been embodied. If the reach-extending baton

affects perception (Witt et al., 2005), why does the automation (which also enhances abilities) not affect perception? Given that the action-specific account of perception would suggest that a change in abilities or performance should lead to a change in perception, it would make sense that the performance boost experienced with automation should lead to a corresponding shift in perception.

So what is the determining factor for why the performance differences caused by automation did not elicit changes in perception in the current experiments? Out of what seems like a complex relationship emerges a simple story when considering motor simulation. As mentioned before, a motor simulation is “the imagining of an action, either covertly or explicitly, without necessarily executing the action” (Witt & Proffitt, 2008). Motor simulation could be a critical factor to explain situations under which the Pong effect, in particular, and action-specific effects, in general, are and are not observed. Indeed, there seem to be many complexities related to when action influences spatial perception, but these complexities can be predicted by appealing to a motor simulation mechanism. The same kinds of situations that evoke a motor simulation are also the kinds of situations that lead to action-specific effects. As will be discussed below, these include anticipating or intending to perform an action, imagining an action, and observing others performing an action (e.g. Grezes & Decety, 2001). Following, I will discuss these effects when watching a computer perform the task, and then will conclude with what this dissertation provides in regards to information about motor simulation when using automation.

It seems that motor simulation is necessary for action-specific effects to emerge. Take, for example, the literature on intention and anticipation. The intention and

anticipation to act leads to a motor simulation of the action that is very similar to what occurs during action execution (Witt & Proffitt, 2008). It could be the result of a motor simulation that provides the information by which perception is influenced by action. To illustrate how this *could* work, consider the action-specific finding that distances appear farther when presented up a hill than on flat ground (e.g. Stefanucci et al., 2005).

Participants could simulate the energy expenditure associated with walking up the hill, and compare that to the energy expenditure associated with walking that same distance on the flat. The comparison between these two motor simulations could be what creates this perceptual distortion.

Appealing to motor simulation could explain why manipulations of effort or ability to perform an action do not always influence perception. Specifically, perception is only influenced by the effort or ability associated with the intended action, so manipulating ability for an action that is not intended has no influence on perception (Witt et al., 2004, 2005, 2010). For example, when the effort it would take to walk a distance increases, perception of the distance to a target increases only for the individual who intends to walk it (Witt et al., 2004). When intending to throw an object to that target, there is no increase in the perceived distance associated with the increase in effort for walking. Conversely, when the effort to throw is increased, the perceived distance to the target is increased as well only when the intention of the actor is to throw but not when the intention is to walk (Witt et al., 2004). This intentionality is likely invoking a motor simulation of the task, which gives rise to information related to the difficulty of the task, which then influences perception. However, when the intended action is different than the action that has been

manipulated, the motor simulation occurs only for the intended action, discounting the irrelevant effort information.

This motor simulation explanation can also explain why results are found even when participants do not have the explicit intention to act nor even the possibility for action. For example, when observers estimated perceived slant, they were never instructed to ascend the hill, yet their ability to ascend the hill related to wearing a backpack or feeling fatigued influenced perceived hill slant (Bhalla & Proffitt, 1999). Similarly, when estimating the distance to targets on a hill and on flat ground, participants were never instructed to walk to the targets, yet they perceived them in relation to the anticipated energy to walk (Stefanucci et al., 2005; Tenhundfeld & Witt, 2017). Moreover, even when action was not possible (such as in virtual environments for which observers were not permitted to move) or when viewing images of staircases, a person's ability to act still influenced perception (Linkenauger, Leyrer, Bülhoff, & Mohler, 2013; Stefanucci et al., 2005; Taylor-Covill & Eves, 2013). Without physical action or even the explicit intention to act, how could action influence perception? The involvement of motor simulation could even explain the emergence of action-specific effects even when there is no explicit intention to act if one were to hypothesize that seeing the ground plane can spontaneously evoke a motor simulation of walking (Witt, Linkenauger, & Wickens, 2016b).

Additionally, motor simulation could account for the host of action-specific effects found when observing an external actor. Bloesch, Davoli, Roth, Brockmole, and Abrams (2012) found that when individuals observed someone else using a reach extending tool, the observer themselves had a perceptual distortion as if they were using the tool. It may seem surprising that perception would be sensitive to another person's actions even when

the observer had no intention to act. But this surprising result can be accommodated by appealing to a motor simulation mechanism. Research has shown that watching someone perform a task elicits, within the observer, a motor simulation of the task as if the observer were acting, possibly as a function of the mirror neuron system (Calvo-Merino, Glaser, Grézes, Passingham & Haggard, 2004; Jeannerod, 1994).

There are two distinct simulation mechanisms that can occur when watching another individual perform a task. One is a mirroring of the other individual's actions. In this instance, it is almost as if the observer becomes the person they are observing (Waytz & Mitchell, 2011). The other form of simulation when observing is as if the observer were to put themselves (and their own abilities) in the situation of the person they are observing. In that instance, it is the observer's own abilities that would matter in the simulation, not those of the individual they are observing. In the first case, the simulation is of the other person's abilities and situation, and in the second case, the simulation is of the observer's own abilities in the other person's situation. In the Pong task, research has shown that the second form of motor simulation appears to occur when watching another human perform the task (Witt, South, & Sugovic, 2014). It is as if the observer puts himself or herself in the position of the person they are observing, and thus their perceptions are distorted in line with the Pong effect.

What about when an observer watches an inanimate object perform a task? The literature is somewhat mixed with respect to the engagement of a motor simulation when watching an inanimate object. For example, motor interference occurs when observing someone else make an incompatible movement but not when watching a robot make an incompatible movement (Kilner, Paulignan, & Blakemore, 2003). This is consistent with

the results of Witt et al. (2012) in which participants who had not yet experienced playing in the Pong task did not show a Pong effect when observing the computer play. Perhaps watching the computer play the Pong task did not invoke a motor simulation. Perhaps they did not put themselves in the shoes of the computer, so to speak, to simulate their own abilities given the situation (which, in this case, refers to the size of the paddle). Given there was no motor simulation, thus, there was no emergent Pong effect (Witt et al., 2012)

In contrast, perception of whether or not a movement is possible is sensitive to movement constraints as described by Fitts' law and this is true for both observing a human and a robot, suggesting no differentiation between the two (Grosjean, Shiffrar, & Knoblich, 2007). This is consistent with the results of Witt et al. (2012) that showed prior experience with the Pong task led to a Pong effect when observing the computer play, even though the computer's ball blocking performance was perfect and thus equal across all paddle sizes. Perhaps having had experience playing the task, participants spontaneously simulated the task as if they were still playing even when watching a computer play.

The mixed results as to when a computer-controlled object invokes a motor simulation (Kilner et al., 2003; Grosjean et al., 2007) can help explain the mixed results found when observing a computer play (as reported in Witt et al., 2012). Perhaps they can also explain the mixed results found in the current experiments.

Given this proposed role of motor simulation for action-specific perception, it is worthwhile to consider the results from the current experiments and the relationship between automation and motor simulation. Automation be likened to the observer-computer conditions for which motor simulation was not evoked, but it could also be considered a tool that improves performance. Thus, on one hand, perhaps when "playing"

with the automated paddle, participants simulate their own abilities as they would be with that paddle had there been no automation. In this case, the results would be similar to if the observer were acting themselves with the various sized paddles. This would be in line with the human results of Witt et al. (2014) who found that an observer puts themselves in the shoes of the other human performing the task. On the other hand, perhaps the observer would simulate their abilities to perform the task with the automated paddle. This is what would be expected if automation were treated like a tool by the perceptual system. This result would be in line with the performance of the automation, as it has become a tool similar to the aforementioned reach extending baton (Witt et al., 2005). The final possibility of motor simulation would be simulating the abilities of the automation. This would be analogous to the aforementioned simulation in which individuals simulate as if they were to take on the skills and abilities of the person they were observing.

So how could motor simulation explain my current results? In Experiment 1, the Pong effect still emerged despite no difference in performance. On the surface this would seem damning for the action-specific account of perception; where there exists no difference in performance, there should exist no difference in perception. However, perhaps the requirement of having participants act at the beginning of each automated trial would have led to a motor simulation of the task. It is possible that the outcomes of the motor simulation (i.e. anticipated performance with the paddle had the participant maintained control) could have outweighed or overridden information related to the automation's performance. This seems analogous to the aforementioned effect where when participants watched a perfect computer perform the task, after having performed the task themselves, there still emerged a Pong effect (Witt et al., 2012). In that case, and in the

trials here, it is likely the case that participants engaged a motor simulation of the task, which would have concerned only their abilities to perform the task, and discarded the performance information of what could be considered an outside actor. This is bolstered by my earlier discussion of a miscalibration to the difficulty. If participants engaged a motor simulation of the task in Experiment 1, it would make sense that the motor simulation would have been more accurately calibrated to the difficulty of the task only after the task had actually been performed. This coincides with what was observed.

In Experiments 2 – 4, there was no effect of either paddle size nor performance on perception. I would argue that this was because when the action was removed, there was no motor simulation. Without a motor simulation, there is no effect on perception. What is different between the automated trials of Experiment 1 and those of Experiments 2 – 4 such that motor simulation was engaged and influential for perception for one but not the others? One option is that in Experiments 2 – 4 there was no action required at all on the automated trials, whereas participants had to position the paddle in order to automation to be activated in Experiment 1. This may have led to there not being any motor simulation when observing the automated trials, which would have left the perceptual system relying on the optically specified information in isolation from the action information.

It is not clear, however, why there would not have been motor simulation on the automated trials in Experiments 2-4 given that I have argued that there would have been in the observational cases of Witt et al. (2012), as discussed above. In both cases the participant would have acted and had experience with the task which would have, theoretically, been sufficient to evoke a motor simulation and subsequent perceptual effect when observing the automated paddle. Perhaps the rapid changes in Experiments 2 – 4

were too quick for the participant to spontaneously engage a motor simulation when observing the inanimate automated paddle. Recall that in Experiment 1, automation was blocked, whereas in Experiments 2-4, automation and controlled trials were intermixed. To be clear, this is distinct from saying that the trials were too quick for motor simulation to be possible. Given that I have argued the Pong effect is perceptual, this would likely have to involve a motor simulation given that the perception of speed would have to occur before the paddle either hit or miss the ball. Otherwise, it would be a judgment effect, which has been ruled out.

It is important to note that it is spontaneous motor simulation that is being discussed here. I could have asked participants to do a task which could have engaged motor simulation (e.g. ask them if they think the ball would be blocked or not), however that was not done. Each of these aforementioned factors could determine if or when spontaneous motor simulations are run. Perhaps spontaneous motor simulation occur when observing a computer if certain conditions are met. Perhaps these conditions include whether the individual has prior experience with the task and whether there is sufficient engagement with the task. This could lead to a motor simulation even when watching an inanimate object perform the task. Future research should determine if these, or other, factors are necessary for this spontaneous motor simulation to occur. Corresponding to these speculations, I contend that automated paddles could influence perception of ball speed but only when they also engage a motor simulation.

### **Framework for a Theory of Automation Embodiment**

The embodiment effects on perception have been well validated in humans as I have described in detail above. Whether it is the inherent ability for one to act (e.g. Taylor et al.,

2011; Witt et al., 2011), performance differences (e.g. Gray 2013; Witt & Dorsch, 2009; Witt & Proffitt, 2005), or even the use of performance enhancing tools (e.g. Davoli et al., 2012; Witt et al., 2005), the claim of action-specific researchers has been that as performance and ability to act change, so too do one's perceptions. This dissertation has shown, however, that it is not as simple as a change in performance leading to a change in perception. When an increase in 'performance' on a task is coupled with the removal of action, as would conceivably be the case in full automation of action control, the perceptual system seems to discard all action information. However, we are left with the question of whether the perceptual system would ever consider automation a tool, much in the same way it considers the paddles tools, and thereby embody it. Said another way, is automation represented the same way by the perceptual system, as is, for example, a stick or the paddles used here.

It is essential, first and foremost, to emphasize that automation is not a single concept. Parasuraman, Sheridan, and Wickens (2000) define automation as "a device or system that accomplishes (partially or fully) a function that was previously, or conceivably could be, carried out (partially or fully) by a human operator." I choose to emphasize this definition because it discusses automation in terms of the human, and not in the abstract. There are two principle components of the definition that deserve parsing out: the use of the term function, and the magnitude qualifiers 'partially or fully'.

In this case function is referring to one of the four distinct functions outline by Parasuraman et al. (2000). Those functions are information acquisition, information analysis, decision and action selection, and action implementation. These functions are the proposed automatable equivalents of the four-stage model of human information

processing of filtering, assessment, decision, and action execution (Onnasch, Wickens & Manzey, 2014; Sebok & Wickens, 2017; Sternberg, 1969). Each of the functions can be automated independently to different levels of automation, which is addressed by the aforementioned 'partially or fully'. While the levels of automation have at times been considered distinct (e.g. Sheridan & Verplank, 1978; Wickens, Mavor, Parasuraman, & McGee, 1998) the levels of automation are now more frequently discussed as a continuum (e.g. Onnasch et al., 2014; Wickens, in press). While the specifics of each level require adaptation to fit each individual function, the anchors of the levels remain the same; at the lowest level the human receives no assistance, whereas at the highest level the human is completely ignored (Parasuraman et al., 2000; Wickens et al., 1998).

The current experiments evaluate the possibility of embodiment at stage four of automation (action support). However, they also speak to automation at stage three (decision and action selection). In Experiment 1, the human chooses (decides) to initiate movement in one direction or the other. The automation then takes over at stage four. Results indicated a residual effect of action embodiment of the paddle sizes. In Experiments 2 – 4, stage three was also automated (in addition to stage four) and this led to no embodiment effect at all. This seems, at least to some degree, to speak to the possibility that preserving human choice, even in the absence of action execution, enables some degree of action automation embodiment. This would be in line with results that suggest that intention to act, alone, can have an affect on perception (Witt et al., 2005).

Beyond the consideration of whether the stage that is automated can have an effect on perceptual embodiment, it is also worth entertaining if the level of automation within each stage can have an effect. There are many examples of partial automation in the real

world. That is to say there are many examples of automation being used that is between the extremes of complete human control, and a fully autonomous system acting independently of the human. Let us take, for example, the bicycle. Because it multiplies power from the input (i.e. the pedaling) we can consider it partial automation of stage four (action execution). This is a lower level of automation than say the motorcycle, which is a lower level still to the full driverless speed and heading control in fully autonomous cars. Because of the increasingly better performance to effort ratio experienced by the individual, the theory of action-specific perception would suggest that as that level of automation increases (thus bettering the performance to effort ratio), there would be a perceptual effect such that distances look shorter and hills look less steep.

In consideration of this gradient, perhaps there is some point at which the automation is no longer seen as inherent to the individual, but is rather seen as an external actor, thereby no longer being embodied. If such a point exists, the initially most obvious place for it would be at the point where the autonomous agent no longer acts in any capacity under the guidance of the human. Whereas at lower levels of automation, the system will still provide information to the human if asked, at higher levels, it will only provide information to the human if it, the autonomous system, decides to, or not at all (respectively) (Parasuraman et al., 2000; Sebok & Wickens, 2017; Wickens, in press; Wickens et al., 1998). At this point one could consider the autonomous system as being completely independent and therefore not under the control, or agency, of the human. However this speculation does necessitate further exploration.

One final consideration to be made regarding the embodiment of automation is the degree to which it is similar to the human. The three features of similarity I think warrant

exploration are behavior, effort expenditure, and performance. Automation behavior is in reference to the degree with which the automation *acts* like the human. There has been considerable exploration into the effects of this type of automation, albeit not in reference to embodiment. For example, research has shown that humans are more likely to trust automation that makes errors similar to those of a human (Madhavan & Wiegmann, 2007), and prefer autopilot systems that controls the airplane like the pilot (Wiener & Curry, 1980). However, what is not clear is whether perceptual embodiment effects elicited by automation that behaved like a human would mean that an individual is actually embodying the automation. I make this distinction in consideration of Witt et al. (2014), which showed that watching another person perform a task results in perceptual embodiment of the observer's abilities rather than those of the person they observe. However, and as mentioned before, this could be partially attributed to the engagement of motor simulation as if the observer was, himself or herself, intending to perform the task (Calvo-Merino et al., 2004; Jeannerod, 1994).

One question for future research aiming to examine the degree to which automation's behavior can lead to perceptual embodiment, is whether it is possible to behave similarly enough to a human to potentially be embodied, without crossing the line which leads to the human simulating their own inherent ability. Additionally, consideration of the automation's behavior should make careful consideration of the consistency between appearance and behavior. There is a well documented phenomena of 'uncanny' characteristics of non-human-like robots leading to the subjective report of the robot seeming 'creepy' or unnerving (e.g. Bethal & Murphy, 2006; Hanson, 2006; Walters, Syrdal, Dautenhahn, Te Boekhorst, & Koay, 2008). While it is unclear whether automation's

behavior that was too human like would decrease the likelihood of embodiment, it is nonetheless worth consideration.

Another potential contributing factor would be the degree to which automation experiences similar effort to the human. In conversations with one of my committee members (C.D.W.), he shared the experience of perceiving a hill to be much steeper when his vehicle's engine was whining trying to get up the hill. What is most interesting about this observation is that this could somewhat be considered almost an effect of empathy on perception. Whereas the driver themselves would not be exerting the extra effort, the embodiment might be of the automation's effort. However, a more concrete explanation in my mind would be the fact that perhaps the whining of the engine is clueing the driver into the difficulty associated with driving up that hill. In the same way that watching someone else perform a task can engage one's brain in a motor simulation of that task (Calvo-Merino et al., 2004; Jeannerod, 1994), perhaps the whining of the engine is somehow engaging a similar motor simulation.

Previous research has shown that a car can be embodied. Individuals who intended to walk a distance perceived that distance to be significantly farther than did those who intended to drive that distance (Moeller et al., 2016). While Moeller et al. (2016) did not evaluate changes in the car's effort, it is clear that the car as a tool can seemingly be perceptually embodied. The key for studies examining automation's effort will be defining what type of effort would lead to an embodiment effect. The action-specific perception literature has traditionally discussed effort in terms of energetic expenditure (e.g. Stefanucci et al., 2005; Tenhundfeld & Witt, 2017; Witt et al., 2004). However, others have shown effects of other cognitive constructs on perceived distance, such as desire (Balcetis

& Dunning, 2010), threat (Cole, Balci, & Dunning, 2013), and social categorization (Xiao & Van Bavel, 2012). For that reason, I believe it to be entirely possible that even though there is no difference in a driver's exertion of energy, the exertion of effort, by an automated system, may be embodied.

The final feature, and the one that can be best addressed by my current data, is the performance of the automation. In addition to allowing me to test whether performance alone was sufficient to drive changes in perception, my experiments utilizing automation represent the first ever examination into whether automation, as a performance enhancing tool, could be perceptually embodied. Given the action specific literature, it seems most likely that if automation were to ever be perceptually embodied, it would be when it performs like a human. The results of Experiments 2 – 4 can speak directly to this.

In Experiments 2 and 3 a small automated paddle was employed that performed as well (and actually better overall) as did the participants with the big controlled paddle. In Experiment 4, the small automated paddle performed the same as did participants with the small controlled paddle. If this performance of automation was to be perceptually embodied by the participants, results should have indicated that the participant's perceptions reflected the overall performance with the automated paddle. However, that was not the case. In all three of those experiments it appears as though there was no embodiment of the automation's performance. With that said, it may be the case that because this form of automation did not *behave* like the human, the fact that it performed like a human may be irrelevant.

However, it is worth noting that these experiments only implemented one slightly odd form of automation. I say odd because automation is almost never switched on and off

every couple of seconds, but is rather chosen to be used for at least minutes at a time. Said another way, the automation used, especially in Experiments 2 – 4, is very untypical of usual automation systems in the real world. It was turned on and off in a random and in a somewhat unpredictable way, every couple of seconds. This is not how automation is typically employed in the real world. Because of this, I would caution extending these null results into saying that automation's performance is not perceptually embodied. Future research could explore whether more real world automation is embodied as well as if being given the choice to utilize automation may lead to embodiment (as was briefly discussed in my theory about the role of agency).

Finally, while each of these possible features of automation have been discussed in isolation, the places in which they intersect may prove to be the most fruitful. While my data only begin to address the perceptual embodiment of automation's performance, the automation in the present set of experiments was designed to not behave like a human. Perhaps that lack of humanlike behavior contributed to the lack of perceptual embodiment.

In summary, this framework for the development of a Theory of Automation Embodiment is still in its early stages. However, given the abundance of literature which details the way in which tool use can affect perceptions, it is worth considering whether automation as a tool can ever be perceptually embodied. Research into this question will need to address the 4 different stages, the 10 different levels, and at least 3 distinct possible features of automation's similarity to a human, as I have only begun to scratch the surface.

### **Applied Considerations**

Having discussed an abundance of theoretical issues, I want to dedicate some thought to more applied considerations, even if just briefly. It is hard to discuss

'automation' without immediately eliciting images of self-driving cars, autopilot, and a whole host of other new-age developments. For that reason, I want to speculate as to what my data could do to potentially inform human factors research. I caution against the reading of any of my speculation as fact, but I do think there exist applied lines of research that could stem from the current data.

The field of action-specific perception as a whole has implications for human factors researchers (Witt, Linkenauger, & Wickens, 2016-a). Given my inclusion of automation, however, I want to focus on where there exists automation of many complex motor tasks such as flying (e.g. Billings, 1997; Parasuraman, Bahri, Deaton, Morrison, & Barnes, 1992) and driving (e.g. Khan, Bacchus, & Erwin, 2012; Stanton & Young, 1998). My data suggest that the automation of these tasks could have potentially devastating consequences. The results of Experiments 2 – 4 suggest that when the ability to act is removed (e.g. through the automation of action/action control), the individual no longer perceives their environment in regards to their ability to act, and thus rely on the optically specified information instead. When we consider that these perceptual embodiment effects are somewhat 'paternalistic' (to steal a phrase from Firestone, 2013), we need to consider what their removal may mean.

Although it has not been shown, let us assume that our perception of speed is directly related to our ability to brake in the car. When our brake pads are worn, or the road conditions are such that it will take longer to brake, our resulting percept should be that we are traveling faster. Research has shown that when we perceive ourselves to be traveling faster, that does in fact change braking behavior (Fajen, 2005). Therefore, a perceptual embodiment of our braking ability would lead us to brake sooner than we

otherwise normally would, to account for our degraded ability to brake. If, however, the action control was automated such that we no longer needed to act, my current results would suggest we might no longer incorporate the action information into the formulation of our percept. Therefore, our decisions to regain control from the automation are based upon a percept that no longer reflects our ability to brake. It is conceivable then that the time at which a driver decides to regain control would be later than necessary in order to safely brake in time to avoid collision.

Similarly, results have shown that performance can also affect perceived runway size (Gray, Navia, & Allsop, 2014). Reliance on perceived size affects the pilot's glide slope, which directly affects their landing performance (Mertens & Lewis, 1981). One could surmise then that automating the control over the aircraft in descent would remove the information about action from the pilot's percept, thereby affecting their perception of runway size as well as of their glide slope. This may lead to a perception of their approach that is not in line with their abilities, which may result in an improper correction following a transfer of control (TOC) from the autopilot to the pilot.

This theory could add to the explanations for why there is such a performance detriment routinely observed in vehicles following a TOC from the automation to the individual (e.g. Merat, Jamson, Lai, Daly, & Carsten, 2014; Samuel, Borowsky, Zilberstein, & Fisher, 2016; Stanton, Young, & McCaulder, 1997). While there exists very compelling explanations like out of the loop unfamiliarity (OOTLUF) (Wickens, 1992), and the lumberjack analogy (e.g. Onnasch et al., 2014), perhaps the action on percepts that are incongruent with one's actual ability to act could explain an additional component of the variance.

While this is entirely speculative, it warrants future exploration into not only whether this discarding of action information exists in flying/driving environments when action control is automated, but whether this could be contributing to observed performance detriments. If it is, consideration into design changes which may elicit embodiment of one's own inherent performance capabilities, even during automated driving/flying (analogous to Experiment 1) may be worth while. However, this would lead to conflict regarding how to ensure the human does not retake control too quickly as they are, for example, perceiving in terms of their very human abilities to react whereas the automated system may be able to react in a much safer and almost 'super human' way. Because of this representing a potentially new frontier in the intersection between perceptual embodiment and the effects of automation, future examination into this could rely on the above framework for a Theory of Automation Embodiment.

## **Conclusion**

By dissociating action capabilities, performance, and paddle size, this dissertation has provided compelling evidence that the Pong effect is not a response bias. This has far-reaching implications for theories of vision; the perceptual system is not modular, but does seem to consider information about action. While this does not prove that all reported top-down effects of cognition on perception are in fact perceptual, this dissertation has helped put one last nail in the coffin to demonstrate that the Pong effect is perceptual. If there is one top-down effect on visual perception, there can be others.

## References

- Balcetis, E., & Dunning, D. (2010). Wishful seeing: More desired objects are seen as closer. *Psychological science*, *21*(1), 147-152.
- Bethal C. L., Murphy R. R.. 2006. Affective Expression in Appearance-Constrained Robots. In Proceedings of ACM SIGCHI/SIGART 2nd Conference on Human Robot Interaction (HRI '06). Salt Lake City, Utah, US, pp. 327-328.
- Bhalla, M., & Proffitt, D. R. (1999). Visual-motor recalibration in geographical slant perception. *Journal of experimental psychology: Human perception and performance*, *25*(4), 1076.
- Billings, C. E. (1997). *Aviation automation: The search for a human-centered approach*.
- Bloesch, E. K., Davoli, C. C., Roth, N., Brockmole, J. R., & Abrams, R. A. (2012). Watch this! Observed tool use affects perceived distance. *Psychonomic Bulletin & Review*, *19*(2), 177-183.
- Calvo-Merino, B., Glaser, D. E., Grèzes, J., Passingham, R. E., & Haggard, P. (2004). Action observation and acquired motor skills: an fMRI study with expert dancers. *Cerebral cortex*, *15*(8), 1243-1249.
- Cole, S., Balcetis, E., & Dunning, D. (2013). Affective signals of threat increase perceived proximity. *Psychological science*, *24*(1), 34-40.
- Cooper, A. D., Sterling, C. P., Bacon, M. P., & Bridgeman, B. (2012). Does action affect perception or memory? *Vision Research*, *62*, 235-240.

- Davoli, C. C., Brockmole, J. R., & Witt, J. K. (2012). Compressing perceived distance with remote tool-use: real, imagined, and remembered. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(1), 80.
- Durgin, F. H., Baird, J. A., Greenburg, M., Russell, R., Shaughnessy, K., & Waymouth, S. (2009). Who is being deceived? The experimental demands of wearing a backpack. *Psychonomic Bulletin & Review*, *16*(5), 964-969.
- Durgin, F. H., DeWald, D., Lechich, S., Li, Z., & Ontiveros, Z. (2011). Action and motivation: Measuring perception or strategies?. *Psychonomic bulletin & review*, *18*(6), 1077-1082.
- Durgin, F. H., Klein, B., Spiegel, A., Strawser, C. J., & Williams, M. (2012). The social psychology of perception experiments: Hills, backpacks, glucose, and the problem of generalizability. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(6), 1582.
- Fajen, B. R. (2005). Calibration, information, and control strategies for braking to avoid a collision. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(3), 480.
- Firestone, C. (2013). How “paternalistic” is spatial perception? Why wearing a heavy backpack doesn’t—and couldn’t—make hills look steeper. *Perspectives on Psychological Science*, *8*(4), 455-473.
- Firestone, C., & Scholl, B. J. (2014). “Top-down” effects where none should be found the El Greco fallacy in perception research. *Psychological science*, *25*(1), 38-46.
- Firestone, C., & Scholl, B. J. (2016-a). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and brain sciences*, *20*, 1-77.

- Firestone, C., & Scholl, B. J. (2016-b). Seeing and thinking: Foundational issues and empirical horizons. *Behavioral and Brain Sciences*, 39.
- Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. MIT press.
- Gray, R. (2013). Being selective at the plate: Processing dependence between perceptual variables relates to hitting goals and performance. *Journal of Experimental Psychology: Human Perception and Performance*, 39(4), 1124.
- Gray, R., Navia, J. A., & Allsop, J. (2014). Action-specific effects in aviation: What determines judged runway size?. *Perception*, 43(2-3), 145-154.
- Grezes, J., & Decety, J. (2001). Functional anatomy of execution, mental simulation, observation, and verb generation of actions: a meta-analysis. *Human brain mapping*, 12(1), 1-19.
- Grosjean, M., Shiffrar, M., & Knoblich, G. (2007). Fitts's law holds for action perception. *Psychological Science*, 18(2), 95-99.
- Hanson D. 2006. Exploring the Aesthetic Range for Humanoid Robots. In Proceedings of Cognitive Science (CogSci 2006) Workshop on Android Science. Vancouver, BC, Canada.
- Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain sciences*, 17(2), 187-202.
- Khan, A. M., Bacchus, A., & Erwin, S. (2012). Policy challenges of increasing automation in driving. *IATSS research*, 35(2), 79-89.
- Kilner, J. M., Paulignan, Y., & Blakemore, S. J. (2003). An interference effect of observed biological movement on action. *Current biology*, 13(6), 522-525.

- King, Z. R., Tenhundfeld, N. L., & Witt, J. K. (2017). What you see and what you are told: an action-specific effect that is unaffected by explicit feedback. *Psychological Research*, 1-13.
- Krueger Jr, N. F. (2003). The cognitive psychology of entrepreneurship. In *Handbook of entrepreneurship research* (pp. 105-140). Springer US.
- Lee, Y., Lee, S., Carello, C., & Turvey, M. T. (2012). An archer's perceived form scale the "hitableness" of archery targets. *Journal of Experimental Psychology: Human Perception and Performance*, 38(5), 1125.
- Leibowitz, H.W. Grade Crossing Accidents and Human Factors Engineering. *American Scientists*, Vol. 95, 1985, pp. 558-562.
- Linkenauger, S. A., Leyrer, M., Bühlhoff, H. H., & Mohler, B. J. (2013). Welcome to wonderland: The influence of the size and shape of a virtual hand on the perceived size and shape of virtual objects. *PloS one*, 8(7), e68594.
- Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of verbal learning and verbal behavior*, 13(5), 585-589.
- Loomis, J. M., Da Silva, J. A., Fujita, N., & Fukusima, S. S. (1992). Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance*, 18(4), 906.
- Madhavan, P., & Wiegmann, D. A. (2007). Similarities and differences between human-human and human-automation trust: an integrative review. *Theoretical Issues in Ergonomics Science*, 8(4), 277-301.

- Merat, N., Jamson, A. H., Lai, F. C., Daly, M., & Carsten, O. M. (2014). Transition to manual: Driver behaviour when resuming control from a highly automated vehicle. *Transportation research part F: traffic psychology and behaviour*, 27, 274-282.
- Mertens, H. W., & Lewis, M. F. (1981). *Effect of different runway size on pilot performance during simulated night landing approaches* (No. FAA-AM-81-6). Federal Aviation Administration, Washington D.C., Office of Aviation Medicine.
- Moeller, B., Zoppke, H., & Frings, C. (2016). What a car does to your perception: Distance evaluations differ from within and outside of a car. *Psychonomic Bulletin & Review*, 23(3), 781-788.
- Onnasch, L., Wickens, C. D., Li, H., & Manzey, D. (2014). Human performance consequences of stages and levels of automation: An integrated meta-analysis. *Human Factors*, 56(3), 476-488.
- Open Science Collaboration. (2012). An open, large-scale, collaborative effort to estimate the reproducibility of psychological science. *Perspectives on Psychological Science*.
- Pagano, C. C., & Isenhower, R. W. (2008). Expectation affects verbal judgments but not reaches to visually perceived egocentric distances. *Psychonomic Bulletin & Review*, 15(2), 437-442.
- Parasuraman, R., Bahri, T., Deaton, J. E., Morrison, J. G., & Barnes, M. (1992). Theory and design of adaptive automation in aviation systems. *Catholic university of America, Washington D.C., Cognitive Science Lab*.

- Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, 30(3), 286-297.
- Proffitt, D. R. (2006). Embodied perception and the economy of action. *Perspectives on Psychological Science*, 1(2), 110-122.
- Poulton, E. C. (1979). Models for biases in judging sensory magnitude. *Psychological Bulletin*, 86(4), 777.
- Samuel, S., Borowsky, A., Zilberstein, S., & Fisher, D. L. (2016). Minimum time to situation awareness in scenarios involving transfer of control from an automated driving suite. *Transportation Research Record: Journal of the Transportation Research Board*, (2602), 115-120.
- Sebok, A., & Wickens, C. D. (2017). Implementing lumberjacks and black swans into model-based tools to support human-automation interaction. *Human factors*, 59(2), 189-203.
- Sheridan, T. B., & Verplank, W. (1978). Human and Computer Control of Undersea Teleoperators. Cambridge, MA: Man-Machine Systems Laboratory, Department of Mechanical Engineering. *MIT*.
- Stanton, N. A., & Young, M. S. (1998). Vehicle automation and driving performance. *Ergonomics*, 41(7), 1014-1028.
- Stanton, N. A., Young, M., & McCaulder, B. (1997). Drive-by-wire: the case of driver workload and reclaiming control with adaptive cruise control. *Safety science*, 27(2), 149-159.

- Stefanucci, J. K., Proffitt, D. R., Banton, T., & Epstein, W. (2005). Distances appear different on hills. *Attention, Perception, & Psychophysics*, *67*(6), 1052-1060.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta psychologica*, *30*, 276-315.
- Sugovic, M., Turk, P., & Witt, J. K. (2016). Perceived distance and obesity: It's what you weigh, not what you think. *Acta psychologica*, *165*, 1-8.
- Taylor-Covill, G. A., & Eves, F. F. (2013). Slant perception for stairs and screens: Effects of sex and fatigue in a laboratory environment. *Perception*, *42*(4), 459-469.
- Taylor, J. E. T., Witt, J. K., & Sugovic, M. (2011). When walls are no longer barriers: Perception of wall height in parkour. *Perception*, *40*(6), 757-760.
- Tenhundfeld, N. L., & Witt, J. K. (2017). Distances on hills look farther than distances on flat ground: Evidence from converging measures. *Attention, Perception, & Psychophysics*, *1-17*.
- Walters, M. L., Syrdal, D. S., Dautenhahn, K., Te Boekhorst, R., & Koay, K. L. (2008). Avoiding the uncanny valley: robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion. *Autonomous Robots*, *24*(2), 159-178.
- Waytz, A., & Mitchell, J. P. (2011). Two mechanisms for simulating other minds: dissociations between mirroring and self-projection. *Current Directions in Psychological Science*, *20*(3), 197-200.
- Wickens C.D. (1992). *Engineering psychology and human performance* (2nd ed.). New York: Harper Collins.

- Wickens, C.D. (in press). Automation stages & levels, 20 years after. *Journal of Cognitive Engineering and Decision Making*.
- Wickens, C. D., Mavor, A. S., Parasuraman, R., & McGee, J. P. (1998). The Future of Air Traffic Control: NRC Report.
- Wiener, E. L., & Curry, R. E. (1980). Flight-deck automation: Promises and problems. *Ergonomics*, 23(10), 995-1011.
- Witt, J. K. (2011-a). Action's effect on perception. *Current Directions in Psychological Science*, 20(3), 201-206.
- Witt, J. K. (2011-b). Tool use influences perceived shape and perceived parallelism, which serve as indirect measures of perceived distance. *Journal of Experimental Psychology: Human Perception and Performance*, 37(4), 1148.
- Witt, J. K. (2017). Action potential influences spatial perception: Evidence for genuine top-down effects on perception. *Psychonomic Bulletin & Review*, 1-23.
- Witt, J. K., & Dorsch, T. E. (2009). Kicking to bigger uprights: Field goal kicking performance influences perceived size. *Perception*, 38(9), 1328-1340.
- Witt, J. K., Linkenauger, S. A., Bakdash, J. Z., & Proffitt, D. R. (2008). Putting to a bigger hole: Golf performance relates to perceived size. *Psychonomic Bulletin & Review*, 15(3), 581-585.
- Witt, J. K., Linkenauger, S. A., & Proffitt, D. R. (2012). Get me out of this slump! Visual illusions improve sports performance. *Psychological Science*, 23(4), 397-399.
- Witt, J. K., Linkenauger, S. A., & Wickens, C. (2016-a). Action-specific effects in perception and their potential applications. *Journal of Applied Research in Memory and Cognition*, 5(1), 69-76.

- Witt, J. K., Linkenauger, S. A., & Wickens, C. D. (2016-b). Action-specific effects in perception and their potential applications: a reply to commentaries. *Journal of Applied Research in Memory and Cognition*, 5(1), 88-93.
- Witt, J. K., & Proffitt, D. R. (2005). See the ball, hit the ball apparent ball size is correlated with batting average. *Psychological Science*, 16(12), 937-938.
- Witt, J. K., & Proffitt, D. R. (2008). Action-specific influences on distance perception: a role for motor simulation. *Journal of experimental psychology: Human perception and performance*, 34(6), 1479.
- Witt, J. K., Proffitt, D. R., & Epstein, W. (2004). Perceiving distance: A role of effort and intent. *Perception*, 33(5), 577-590.
- Witt, J. K., Proffitt, D. R., & Epstein, W. (2005). Tool use affects perceived distance, but only when you intend to use it. *Journal of experimental psychology: Human perception and performance*, 31(5), 880.
- Witt, J. K., Schuck, D. M., & Taylor, J. E. T. (2011). Action-specific effects underwater. *Perception*, 40(5), 530-537.
- Witt, J. K., South, S. C., & Sugovic, M. (2014). A perceiver's own abilities influence perception, even when observing others. *Psychonomic bulletin & review*, 21(2), 384-389.
- Witt, J. K., & Sugovic, M. (2010). Performance and ease influence perceived speed. *Perception*, 39(10), 1341-1353.
- Witt, J. K., & Sugovic, M. (2012). Does ease to block a ball affect perceived ball speed? Examination of alternative hypotheses. *Journal of Experimental Psychology: Human Perception and Performance*, 38(5), 1202.

- Witt, J. K., Sugovic, M., & Taylor, J. E. T. (2012). Action-specific effects in a social context: others' abilities influence perceived speed. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(3), 715.
- Witt, J. K., Sugovic, M., Tenhundfeld, N. L., & King, Z. R. (2016). An action-specific effect on perception that avoids all pitfalls. *Behavioral and Brain Sciences*, *39*.
- Witt, J.K., Tenhundfeld, N.L., & Tymoski, M. (in press). Is there a chastity belt on perception? *Psychological Science*.
- Wraga, M., Thompson, W. L., Alpert, N. M., & Kosslyn, S. M. (2003). Implicit transfer of motor strategies in mental rotation. *Brain and Cognition*, *52*(2), 135-143.
- Xiao, Y. J., & Van Bavel, J. J. (2012). See your friends close and your enemies closer: Social identity and identity threat shape the representation of physical distance. *Personality and Social Psychology Bulletin*, *38*(7), 959-972.

## Appendix A

*PSEs and 95% Confidence Intervals in centimeters per second for the No-Automation Start condition in Experiment 1*

	<u>Automated</u>			<u>Not Automated</u>		
	Small	Medium	Big	Small	Medium	Big
PSE (cm/s)	43.56	46.18	47.88	43.61	45.98	49.45
95% CI (Within)	2.35	1.96	2.35	2.44	2.43	1.87

*PSEs and 95% Confidence Intervals in centimeters per second for the Automation Start condition in Experiment 1*

	<u>Automated</u>			<u>Not Automated</u>		
	Small	Medium	Big	Small	Medium	Big
PSE (cm/s)	44.65	44.99	47.61	43.25	46.36	49.81
95% CI (Within)	1.50	1.35	1.66	2.00	1.50	2.15

*PSEs and 95% Confidence Intervals in centimeters per second for Experiment 2*

	Small Controlled	Small Automated	Big Controlled
PSE (cm/s)	42.69	45.63	47.49
95% CI (Within)	0.93	1.22	0.83

*PSEs and 95% Confidence Intervals in centimeters per second for Experiment 3*

	Small Controlled	Small Automated	Big Controlled
PSE (cm/s)	44.70	46.66	47.78
95% CI (Within)	0.64	1.05	0.71

*PSEs and 95% Confidence Intervals in centimeters per second for Experiment 4*

	Small Controlled	Small Automated	Big Controlled
PSE (cm/s)	43.21	44.96	47.32
95% CI (Within)	0.53	0.94	0.81