The Devil is in the Distribution:
Refining an ACT-R model of a Continuous Motor Task

Melissa A. Gallagher (melissa.gallagher@rice.edu)
Department of Psychology, MS-25
Houston, TX 77005 USA

Michael D. Byrne (byrne@rice.edu)
Departments of Psychology and Computer Science, MS-25
Houston, TX 77005 USA

Abstract

A continuous motor task was reexamined to increase the fidelity of a previous model of the task. Previous modeling work was successful in matching qualitative performance, but not all aspects of the movement profile. When a new set of subjects were brought into the lab and eye-tracked, their motion data informed further modifications to the model's performance strategy as well as changes to the ACT-R architecture. This in turn produced a model with a higher fidelity movement profile. The change in strategy modeled identified capabilities that could be implemented in future work as well as brought new inspiration to strategy refinement on this task.

Keywords: Human motor movement; motor control; continuous motor task.

Introduction

Manual control and tracking tasks have been well researched in terms of two models, the "crossover" model and the optimal control theory model (Jagacinski & Flach, 2003). These models have adequately explained human manual tracking in many contexts. In the age of automation, these tasks are becoming less frequent, but the ones that still occur are very important. For example, as new surgical tools are developed, new training regimens will have to be created for surgeons to learn how to use them. Inadequate training can have high costs, such as injury or loss of life. An additional motivating factor for creating good training regimens is that the amount of time it takes to learn a task is decreased, producing long term cost savings. Many of these tasks have interesting cognitive components as well as motor components, and the intersection of these components make it a rich research area. Cognitive models are especially useful in situations where there is a high cost for error and all of the components that comprise the optimal strategy for the task are not completely understood.

The Task

For our efforts we looked at a particular task that had been studied extensively by O'Malley and colleagues (O'Malley, et al., 2006; Li, Patoglu, & O'Malley, 2009; Huegel, et al., 2009; Powell, 2010). The task layout is depicted in Figure 1. This task is a challenging motor control task but it is not a tracking task because the targets are stationary and the user generates the only movement. The task is scored on how many targets the subject can hit in a 20 second interval.

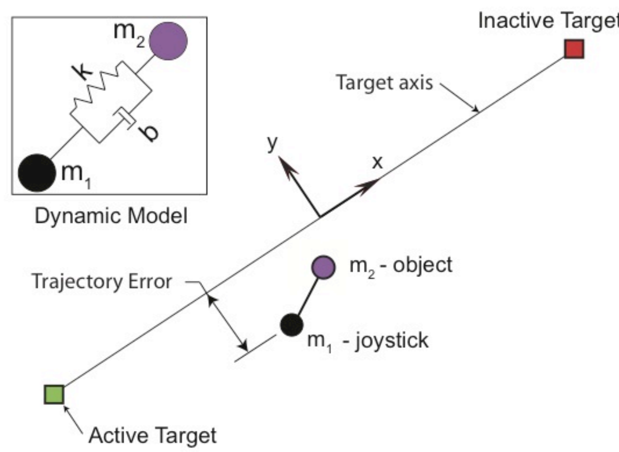

Figure 1: The task configuration. m_1 is controlled by the joystick; the goal is to alternate hitting the two targets with m_2. m_2 is attached to m_1 by a spring with a spring constant k and a damping coefficient of b.

The work done by O'Malley et al. initially focused on different training methods to improve the subjects' performance on this task. All of the subjects were instructed that the optimal strategy for task performance was to move between two points on the screen that were symmetrical across the y-axis, at rate that corresponded to the natural frequency of the system. Symmetrical movements travel approximately the same distance on the target axis on either side of the y-axis. The stiffness and damping parameters determine the natural frequency of a system and in this case created a slightly underdamped system. When an underdamped system is excited at or near its natural frequency the output amplitude will be greater than the input amplitude. This allows the subjects to get a higher hit count while inputting less energy into the system. Subjects were then separated into groups to determine if different training methods affected their performance on the task. Some of the training methods that were employed were haptic feedback and visual cues. The results of the training studies did not support the use of any one particular training method to improve performance. However, the studies showed an interesting pattern of subjects' learning over the sessions.

When the data of the subjects' hit performance was analyzed, three distinct groups appeared. The first group was the low performers, which were classified as such if their final score fell below one standard deviation from the mean. The subjects in this group started out with poor performance on the task and only made modest linear improvements over the experimental sessions. The movement profile for the low performers suggested that they tended to make a circular motion around the field, effectively swinging the coupled mass around to hit the targets. The second group of subjects was the high performers, which were classified as such if their initial performance was one standard deviation above the mean. The subjects in this group started out with good performance on the task and showed the same pattern of linear improvement over sessions. The movement profile for these subjects suggests that they moved along the target axis with minimal off-axis motion, exploiting the physics of the system to improve their performance to get over one target hit per second. The third group of subjects made up a group called the transitional group. This group started out with performance comparable to the low performers but ended up with performance comparable to the high performers. The learning curve that the transitional performers exhibited was better fit by a logarithmic function than a linear one. The movement profiles for the transitional group started similarly to the low performers and end up similarly to the high performers. What information the transitional subjects learned to change their performance from low to high is currently an open question. The three performance groups naturally formed in spite of the fact that all subjects were given the same instructions for optimal performance on the task. Figure 2 displays the performance of the three groups over the sessions.

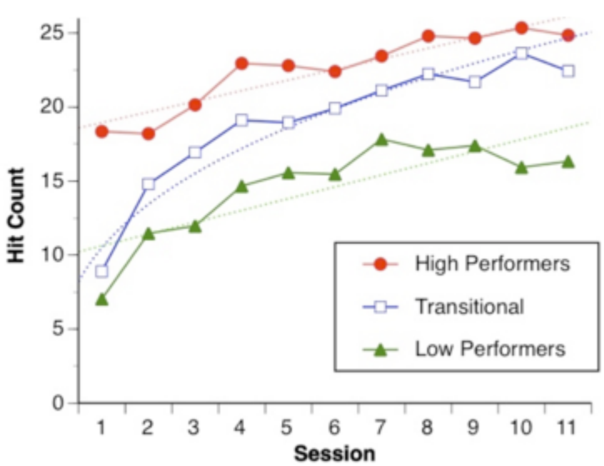

Figure 2. The mean hit count for the three different subject groups over session. Session 11 was a retention session.

Previous Modeling Work

The research done by Huegel (2009) included an in-depth analysis of the movement profile of the high performers. This allowed the model to be evaluated not only on the hit count, but also on whether or not the performance strategies matched. A Fourier analysis of the high performers' movement profile showed that they moved at a consistent frequency that was either at or slightly higher than the natural frequency of the system. This allowed the high performers to have a higher hit count without losing control of the system. Huegel's analysis also found that the high performers' motion was regular, symmetrical, and almost entirely on axis. The on-axis movement and highly regular motion explained most of the variance in the number of target hits between individuals and between sessions. These two variables were only weakly correlated with each other and they appear to be separable components of the task. Additionally, the high performers also made an initial ramp-up movement to excite the system and get a target hit on the first motion, increasing their hit count across trials.

A number of modifications were made to the ACT-R architecture so that the expert performance on the spring-target task could be modeled. The modifications made were to the manner ACT-R moved the mouse, the output of the aimed movement, and how ACT-R handles movement to targets not drawn on the screen (Byrne et. al. 2010). The minimum jerk velocity profile was implemented because of past success modeling reaching and aimed movement with it (Hogan, 1984). With these modifications to ACT-R in place, a model of the central tendency of the high performers was constructed. This model utilizes the same strategy that Huegel's high performers were instructed to use. The model regularly oscillates between two points along the target axis. These two points were called virtual targets and they are represented by squares of a specific size with a constant center. The virtual targets were positioned along the target axis and symmetrically across the y-axis and so that the movement between the two targets was at the natural frequency of the system. The virtual targets are not visually represented on the screen and make use of the imaginal buffer in ACT-R. Visual attention shifted between these two points with the movement even though there was no visual representation of the target. Over the course of 200 runs of the model, a mean hit count of 23.96 hits per 20-second trial was achieved. The largest mean for the high performers in sessions 10 was 25.35. There is a 5% difference between the model's performance and the human performance. In addition to the lower hit count this model has some discrepancies in performance compared to the high performers in Huegel's study, mainly in end point noise. ACT-R currently uses a Gaussian distribution for on-axis and off-axis error while the subjects have smaller off-axis error than the model would predict.

After creating a model that can perform as well as the high performers at the end of the task, efforts were focused on creating a model that would refine its strategy and produce the same learning pattern as the high performers. The learning models used the same strategy of oscillating between virtual targets but employed different means of refining the strategy to improve performance. The first method of strategy refinement was to have the model learn the size of the virtual target. The model would start with a random size for the virtual target that generated stable performance and would increase the size, decrease the size, or keep the size the same. The second method of strategy refinement was to give the model a set of competing virtual target locations and have it learn which target locations were the optimal ones to move to. For the learning models, a slightly lower performance would be expected because the ramp up function was excluded, but neither of these models was able to achieve the same strategy refinement over the simulated sessions as the high performers. The model that learned the size of the virtual target would occasionally generate good performance for a trial but over the simulated sessions would not converge to a stable target size. The model that learned the locations of the virtual targets started with performance that was similar to the high performing subjects during their first session but over the simulated sessions it was unable to refine the strategy and performance never improved.

The Experiment

Due to the fact that neither of these learning strategies produced evidence of the model refining its strategy, the motion data was reexamined. The strategy being modeled was reevaluated as some evidence suggested that even though the subjects were given a particular strategy to use, they might be using a different strategy, such as monitoring the physics of the system rather than moving between virtual locations. A follow-up study, which eye-tracked high performing subjects, was run to gain insight into the task strategy they might be using. The patterns in the eye data were used to help determine whether the strategy might be similar to the model, which oscillates its visual attention between virtual targets, or an alternate strategy of monitoring the physics of the system.

Method

Subjects. Eleven subjects were identified as high performers from their performance in the final session of Powell's study (2010). Five of the subjects participated in the eye-tracking study. All subjects were male, right handed and between 19-24 years old. The subjects received ten dollars for their participation.

Apparatus. The experiment was run using a 1.8 GHz Macintosh G5 PowerPC running MAC OS 10.3.9. Stimuli for the experiment were displayed on a 19" Sony CRT monitor at a frame rate of 60 hz and a resolution of 1152 by 864 pixels. Subjects were seated directly in front of the display and interacted using a mouse. The same machine was used for all participants. The eye tracker used was an ISCAN RK726/RK520 HighRes Pupil/CR tracker with a polhemus FASTRACK head tracker. Head-mounted optics and a sampling rate of 60 Hz were used in the experiment.

Stimulus and Design. The stimulus was the spring target task described previously. Subjects were directed to get the most hits possible during each trial. Each subject did five practice trials and then 20 trials of the task. The practice trials were to get accustomed to the eye-tracker and mouse and for the subjects to reacquaint themselves with the task. Hit count, motion data, and video of the eye movements were collected during all trials. The practice trials were excluded from analysis.

Procedure. When subjects arrived they were reminded of the task procedure and told to perform the task as they had done previously. They were then calibrated with the eye-tracker and after a successful calibration the subject began the practice trials. After the task was over the subjects were debriefed and paid for their time.

Results

The mean hit count of the subjects in the eye-tracking session was 19.78 with a standard deviation of 2.02 and a range from 6 to 27. In addition to the hit count the movement profiles of the subjects were examined. To analyze the movement profiles of the high performing subjects seven metrics were calculated for each trial. These metrics were amplitude, peak-to-peak amplitude, number of movements, trajectory error, total distance traveled, pause time between movements, and movement time. A correlational analysis showed that only two of the seven performance metrics reliably correlated with hit count. The first metric is the trajectory error (r = -.34, p = .002), measured as the number of pixels deviated from the target axis. The second variable is the peak-to-peak amplitude (r = .28, p = .005), measured as the distance traveled for each movement. The subjects employed a variety of visual strategies for monitoring the objects on the screen. Among these were oscillating between the targets, keeping the eyes fixed in the center of the screen, and oscillating between two points on the screen between the targets. In 46% of the trials recorded the subjects employed more than one visual strategy during the trial. To analyze the eye data each trial was classified according to which strategies the subject used during that trial by watching the playback. The visual strategy did not have a reliable relationship with hit count. Since the movement is harmonic, the peak velocity of the subject's on axis movement was also evaluated to inform further design decisions. It was found to be higher than the model's. The graph of on-axis velocity over time was examined for all subjects. For one trial, the subject's mean peak velocity was 0.58 pixels per millisecond while the model's mean peak velocity was 0.45 pixels per millisecond.

Experiment Discussion

Since the eye tracking session occurred six months after the subjects were trained on the task, the drop in performance was not unexpected as even a month after training the performance drops. An additional factor in the performance drop could be that, when the subjects were trained on the task they used a high quality haptic joystick and during the eye tracking session they used a computer mouse. This was due to the fact that we did not have access to the joystick and the ACT-R model uses a mouse. Even with these differences in environment, differences between the model's behavior and the subject's behavior became evident during the analysis.

Since the peak velocity of the model movement is bound by Fitts' Law and was lower than the subjects', to have the model move faster the virtual target size needs to increase, the distance between the virtual targets needs to decrease, or both. With a faster velocity the model will be able to make more movements and get more hits during the twenty seconds of the trial. To be able to increase the virtual target size a better model of the end point error distribution is needed to not introduce more off-axis error, which is moderately correlated with hit count. Additionally the relationship between the peak-to-peak amplitude and hit count indicated that the model needed to be changed so that it does not make symmetrical movements across the midpoint of the target axis. As non-symmetrical movements are not bound to a single point across the axis, the model has an increased number of virtual targets it is able to move to. This makes it possible for the model to produce movements of more consistent peak-to-peak amplitude and in some cases shorter movements. The ability to make movements more frequently will contribute to a higher hit count.

Modeling

The two metrics that had a reliable relationship with hit count informed the efforts to modify the model to be more in line with the central tendency of the subjects' behavior. The following describes modifications to how ACT-R handles the endpoint error model of a movement and how the model's behavior was changed to make non-symmetrical movements across the axis.

End Point Error Model

When a target is clicked on, very rarely is it clicked on in the exact center. May (2012) studied where the final landing point is for symmetrical two-dimensional targets. The model separates endpoint error into two components, on-axis and off-axis error. The on-axis error is measured by how far from the center of the target the end point is on the axis of approach. The off-axis error is measured as the distance from the center of the target on the axis that is perpendicular to the axis of approach. May found that the end point error distribution is wider for on-axis movement than for off-axis movement. The on-axis end point error is distributed normally with a 96% accuracy rate across the effective width of the distribution. The off-axis end point error has the same distribution but is scaled to be over only 75% of the off-axis target width.

The current implementation of end point error in ACT-R does not distinguish between on-axis and off-axis error. The method samples from a logistic distribution where 96% of it is half the target width. The randomly sampled value becomes the magnitude of the error. A vector of this magnitude is then added to the center of the error in a random direction. Modifying how noise is added to the ending position of the mouse movement to be more in line with the model of May is advantageous to modeling this task as large off-axis error is detrimental to achieving a high hit count on the task. Since the high performers have low off-axis error this model of end point error distribution allows the model to make quicker movements because the target width can be increased without the model's performance suffering from more off-axis error.

The model of expert behavior was modified to use virtual targets of a size of 19 pixels increased from the previous size of 14 pixels. This size increase was selected so that 75% of the off-axis width of the new virtual targets was approximately 96% of the off-axis width of the previous virtual target size used. This modification allows the model to make faster movements, as the target is larger, while not increase the amount of off-axis error and reducing hit count because the effective error size is the same when using the different end point error distribution.

Non-Symmetrical Movement

Non-symmetrical movements means that the distance traveled on either side of the y-axis does not have to be the same. To create non-symmetrical movements in the model a set of fourteen virtual targets were defined. Half of the virtual targets were on the left side of the midpoint of the target axis and the other half were on the right side. The virtual targets covered the majority of the range of end points that the subjects generated during the experiment. The area of the virtual targets overlapped so that there were no areas along the target axis within this range that had an extremely low probability of being selected. Figure 3 shows the subject generated endpoints in red and the virtual target locations in blue. Since these virtual targets covered a wide area a matching system was employed to keep the peak-to-peak amplitude consistent between movements. The matching system ensured that only productions that were within the range of the optimal amplitude could be selected sequentially. For most virtual targets there were three possible virtual targets on the opposite side of the off-axis that could be the destination for the next movement. Each of these three virtual targets had an equal likelihood of being selected. Once a virtual target was selected for the next movement a new set of three virtual targets would be chosen from for the movement following. After the ramp up movement, all seven productions were equally likely to fire and only for movements after that was virtual target location restricted by the amplitude to the next virtual target.

Results

The movements of the model were made at a rate of approximately 1.4 Hz. The noise in the model was generated from both the end point distribution and the different time duration of perceptual motor operations. For the primitive operators, the standard noise setting in ACT-R was used with no parameter tuning. Over 200 runs of the model it produces a mean hit count of 24.93 hits per 20-second trial with a standard deviation of 1.25 and a range from 22 to 27. For one run of the model the new mean peak velocity was 0.65 pixels per millisecond.

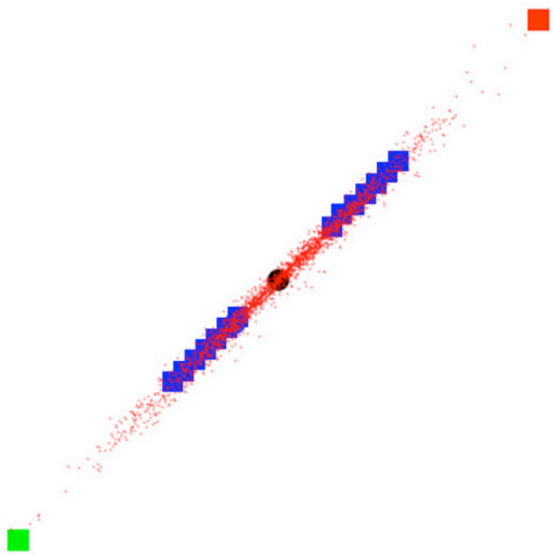

Figure 3. The task layout with the movement endpoints and virtual target locations overlaid. The endpoints the subjects generated are shown in red and the area the virtual targets covered are outlined in blue.

Discussion

This model demonstrates that with additional capabilities in ACT-R's motor module it is possible to model expert performance on this task after training. While there is only a small increase in hit count with these modifications to the original model there is greater consistency with the subject's motion data. One of the first important characteristics of the new model's movement is that it moves at a velocity that is similar to the subjects'. There is also greater variation in the velocity of the new model as there is more variation in the distance that it travels. Furthermore the trajectory of the new model has more noise than the previous model though not quite as much as the subjects (Figure 4). Modeling the correct trajectory and velocity profile is essential for gaining further understanding into task strategy and refinement.

In continuing work, efforts to model learning on this task can be undertaken with a fresh perspective on the expert performer's strategy. The previous efforts of attempting to model learning with a set of competing virtual targets had the targets bound symmetrically across the off-axis for a set of two movements. While the model produced no evidence of strategy refinement with this method, it generated performance that was stable and produced a hit count similar to the high performers during the first session.

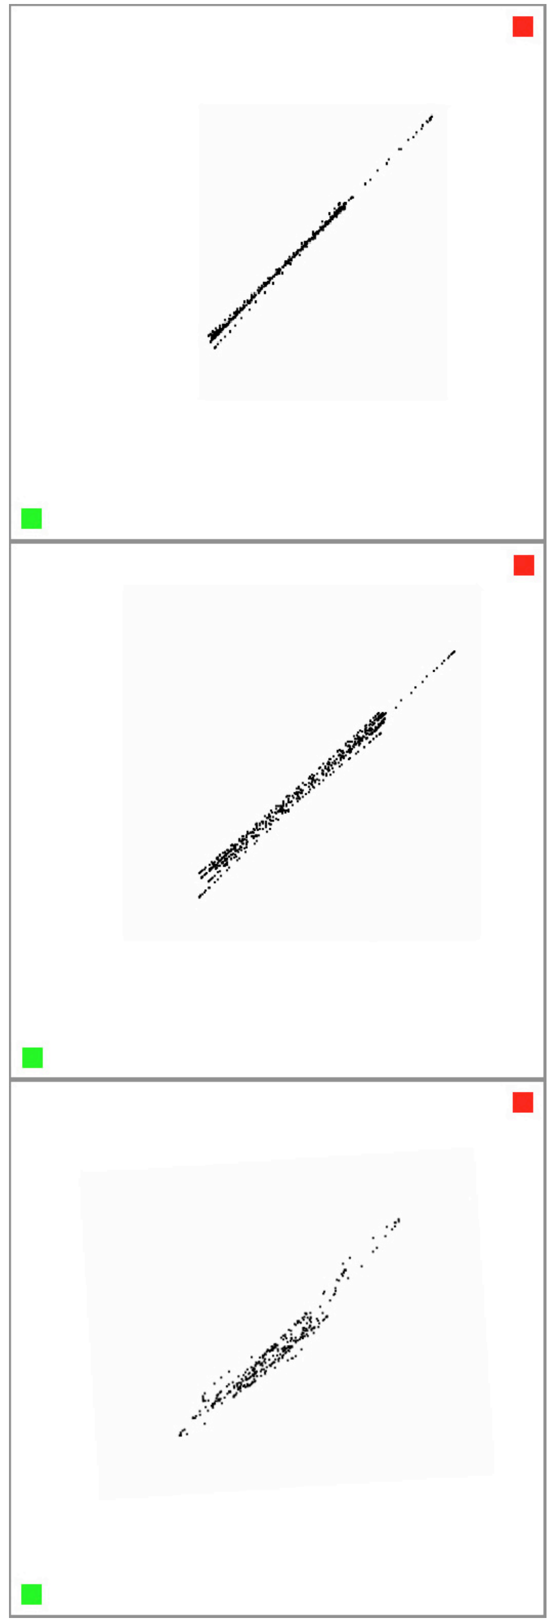

Figure 4. Three trajectories for trials with 22 hits. The top trajectory was generated by the previous model, the middle by the new model, and the bottom is one subject's data.

With the restriction of symmetrical movements lifted, a new method of strategy refinement can be tested. This strategy refinement could be learning the optimal movement amplitude, which would be represented by which pairs of virtual targets should be associated with each other rather than the strategy of which virtual targets should be moved to during the trial.

While improved performance of modeling expert behavior on this task brings hope to modeling expert strategy refinement, this task highlights other types of movement that could be implemented in ACT-R to have a more complete model of human motor movement.

Modeling low performers' behavior on this task is not currently possible in ACT-R as low performers make curved movements rather than linear movements. While the Steering Law (Accot & Zhai, 1997) generates accurate predictions for the movement time along a curved path, there are competing theories as to the velocity profile that humans use to generate this type of movement. Modeling the strategy that the low performers use is also more difficult than modeling the high performers strategy. The high performers strategy of oscillating on-axis movement with little attention being paid to the system only required a few extensions to the architecture to implement in ACT-R. The low performers strategy of swinging the coupled mass around the field requires the model to pay attention to where the disc is on the screen in relation to the active target and for the model to have some knowledge of how the movement of the tool will affect the position of the mass. Not only are these movements not aimed, it is also unclear how the subject plans the movement path. Additional work into curved movement and path planning needs to be done to be able to extend ACT-R to have the capability to model the low performers on this task.

The minimum jerk profile performs well on this task and has found success in describing human behavior on reaching tasks; it is limited in describing other types of movement. Another drawback is that it is smoother than the movements that humans actually make. For two separate movements that travel the same distance in the same amount of time it will generate the same velocity profile for both movements, which human subjects will rarely do. For more motor tasks to be modeled using ACT-R other velocity profiles will need to be implemented, especially in physics-based tasks like this one. A ballistic movement profile could also model the subject's movement due to its similar shape in velocity profile but markedly different one on acceleration. As modeling extends beyond the point-and-click paradigm to tasks where continuous motion matters, different velocity profiles will need to be implemented to give a better understanding of task performance. There is rich potential for cognitive modeling to grow in the motor domain, not only in current capabilities but also in learning procedures, path planning, and other more difficult skills.

Acknowledgments

This research was supported by the National Science Foundation under grant #IIS-0812569. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied, of the NSF, the U.S. Government, or any other organization.

References

Accot, J., & Zhai, S. (1997). Beyond Fitts' Law: Models for trajectory-based HCI tasks. *CHI '97 Proceedings of the SIGCHI conference on Human factors in computing systems*, 295-392.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.

Byrne, M. D., O'Malley, M. K., Gallagher, M. A., Purkayastha, S. H., Howie, N., & Huegel, J. C. (2010). A Preliminary ACT-R model of a continuous motor task. In *Proceedings of the Human Factors and Ergonomics Society 53nd Annual Meeting* (pp. 1037-1041). Santa Monica, CA: Human Factors and Ergonomics Society.

Hogan, N. (1984). An organizing principle for a class of voluntary movements. *Journal of Neuroscience*, 4, 2745–2754.

Huegel, J. (2009). Progressive Haptic Guidance for a Dynamic Task in a Virtual Training Environment. Doctoral Dissertation, Rice University, Houston, TX.

Huegel, J., Celik, O., Israr, A., & O'Malley, M. K. (2009). Expertise- based performance measures in a virtual training environment. *Presence: Teleoperators and Virtual Environments, 18*(6), 449– 467.

Jagacinski, R. J., & Flach, J. M. (2003*). Control theory for humans: Quantitative approaches to modeling performance.* Mahwah, NJ: Erlbaum

Li, Y., Huegel, J., Patoglu, V., & O'Malley, M. K. (2009). Progressive shared control for training in virtual environments. In *Proceedings of the International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems and the Third Joint World Haptics Conference (HAPTICS)* (pp. 332–339).

Li, Y., Patoglu, V., & O'Malley, M. K. (2009). Negative efficacy of fixed gain error reducing shared control for training in virtual environments. *ACM Transactions on Applied Perception, 6, 1–21.*

May, K. (2012). *A model of error in 2D pointing tasks.* Undergraduate Honors Thesis, Rice University, Houston, TX.

O'Malley, M. K., Gupta, A., Gen, M., & Li, Y. (2006). Shared control in haptic systems for performance enhancement and training. *ASME Journal of Dynamic Systems, Measurement and Control, 128, 75–85*.

Powell, D. (2010). Implementation and Analysis of Shared-Control guidance Paradigms for Improved Robot-Mediated Training. Masters Thesis, Rice University, Houston, TX.