Transfer effects of varied practice and adaptation to changes in complex skill acquisition

Roderick Yang Terng Seow (yseow@andrew.cmu.edu)
Shawn Betts (sabetts@andrew.cmu.edu)
John R. Anderson (ja0s@andrew.cmu.edu)
Department of Psychology, Carnegie Mellon University, 5000 Forbes Ave
Pittsburgh, PA 15213 USA

Abstract

Varied training in comparison to consistent training has been shown to benefit transfer to novel conditions within the motor learning paradigm. However, it is unclear if these benefits of variable training extend to complex skills such as driving. Unlike simple motor skills, these complex skills require individuals simultaneously to learn the mapping between ones actions and their consequences and also to integrate this knowledge into continuous and dynamic responses to the changing demands of the environment. In the current work, we compare observed data and an ACT-R model of complex skill acquisition on a navigational video game task (Space Track). Participants trained either on one or two levels of thrust. Performance on a transfer test was better in the varied training conditions in both humans and model. Performance in both humans and model was also differentially influenced by the most recently practiced thrust level. Further analyses revealed large differences between model and human behavior on more detailed measures, which suggests that that the model achieves the same overall performance through different strategies. We discuss these findings and their implications for the ACT-R model of skill acquisition.

Keywords: varied practice; transfer; adaptation; ACT-R; complex skill acquisition

Introduction

Transfer learning is the phenomenon in which practice on one task facilitates the learning of a related task, which reduces the time needed to attain a certain level of skill. This phenomenon has been studied in various domains such as mathematical problem solving (e.g. Speelman and Kirsner, 2001), perceptual categorization and discrimination (e.g. McGovern et al., 2012), and sensorimotor learning (e.g. Goodwin et al., 1998). Within the domain of sensorimotor learning, practicing on varied task parameters has been shown to facilitate more transfer to new task parameters as compared to consistent practice. For instance, when the transfer task is to toss a beanbag to a target at a set distance, participants who trained on different target distances excluding the transfer target perform better than those who trained on just one target distance (Kerr and Booth, 1978).

In the tasks used to investigate the effects of varied practice, one common feature is that the goal of the task is often closely related to the sensory consequence of the sensorimotor mapping that the subject needs to learn. For example, the goal in a visuomotor rotation task is to maneuver a cursor towards a virtual target in the presence of perturbations (e.g. Braun et al., 2009). These perturbations cause the motion of the cursor to rotate with respect to the motion of the controller and successful participants are hypothesized to learn this new mapping between the motion of their hand and the motion of the cursor. However, it remains unclear is unclear if the benefits of varied practice extend to more complex tasks in which the acquisition of sensorimotor maps is necessary but insufficient for high performance. For example, assuming that the goal of a driver is to get from point A to B in the fastest and safest manner, successful driving involves not only more than just learning how the movement of ones foot on the accelerator translates to the cars motion, but also the ability to come up with an action plan to navigate the upcoming obstacles or road hazards. Hence, one of the goals of this study is to answer the following question: When learning complex skills, where the sensorimotor map is only a part of the skills needed to accomplish the task goal, does varied training still outperform consistent training with regard to the transfer of performance to novel task parameters?

Task

Space Track was originally a video game developed by Anderson et al. (in press) as part of a study on the transfer of complex skills. Just like driving, mastering Space Track is a complex skill because it requires one to integrate perceptual, motor, and cognitive components. Expertise arises from having gained an intuitive understanding of the physics of the game and the ability to use that knowledge towards planning sequences of key presses to overcome various situations. In Space Track, players control a space ship in a frictionless environment using three keys: thrust (W), rotate clockwise (A), and rotate counterclockwise (D). Players earn 25 points by successfully navigating the ship along each rectangular track segment and lose 100 points when the space ship crashes into the walls of the track. Figure 1 shows a schematic of the task. Finding a good speed is crucial for performance one needs to fly fast enough to cover as much distance as possible but also slow enough to avoid losing control of the ship and crashing.

To create changes in the task environment, we manipulated the amount of thrust the ship receives for the same duration that the thrust key is depressed. When the thrust key is depressed, a vector of x pixels per second in the current direction of the ship is added each game tick, which is 1/30th of a second. For the same duration of key press, a game with higher thrust would cause the ship to fly faster than a game with lower thrust. Mastery of the game relies on adequately predicting and controlling the motion of the space ship. Thus, players would have to retune their control parameters when faced with a different thrust level.

We created three game types, each with a different thrust level. High thrust games (H) added 0.6 pixels / tick to the ships velocity vector for each tick that the thrust key was depressed. Medium (M) and low (L) thrust games added 0.4 and 0.2 pixels / tick respectively. With these three game types, we created four training conditions as follows: LLLLM, HHHHM, LHLHM, and HLHLM, where each letter stands for one block of 8 x 3-minute games. Figure 2 provides a pictorial representation of the task design. For instance, a player in the LLLLM condition would play 4 blocks (32 games) of low thrust followed by 1 block (8 games) of medium thrust. For our analyses, the first 4 blocks will be referred to as the training blocks, and the last block of medium thrust in all conditions will be referred to as the test block. Participants in the consistent training group will be assigned to either LLLLM or HHHHM, while those in the varied training group will be assigned to either HLHLM or LHLHM. Our rationale for using two different conditions in the consistent training group is to separate adaptation effects due to consistent vs. varied training and effects due to training on a high vs. low thrust. We used two different conditions in the varied training group to account for possible block order effects.

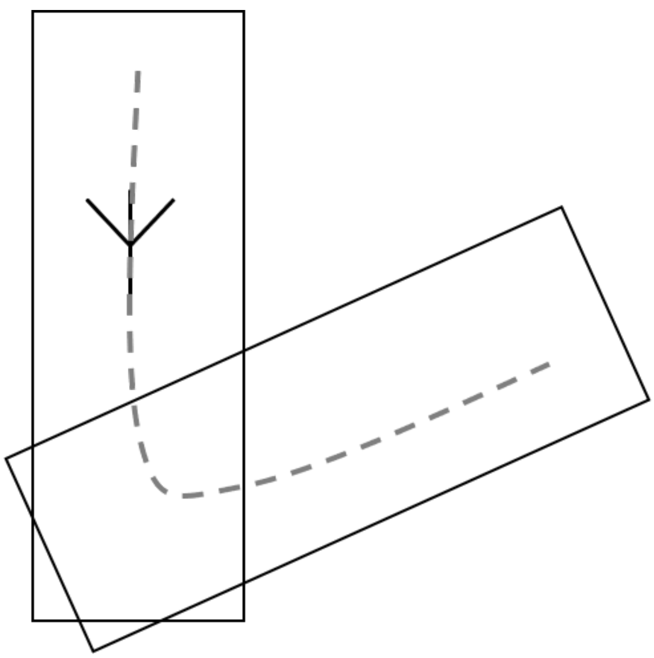

Figure 1: Schematic of Space Track. The goal is to navigate a space ship along a racetrack with rectangular segments. The dashed line displays a potential trajectory along two consecutive segments.

Experiment

80 participants, 22 females and 58 males, ranging in age from 21 to 65 years old (mean = 31.0) completed both experiment sessions through Amazon Mechanical Turk. Participants were paid $5 for the first session and $10 for the second session plus a bonus of $0.03 per 100 points.

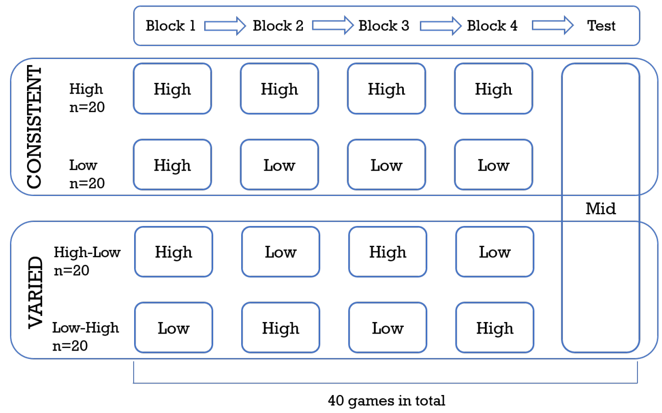

Figure 2: Task design. Each row represents one condition, and each box represents one block of 8 games.

The experiment consisted of two sessions. In the first session, participants filled out a demographic questionnaire, then proceeded to complete the first 20 games. Participants that passed a set of inclusion criteria were then invited to the second session, which consisted of another 20 games. During a game, if 20 seconds elapse without the participant pressing a key, a pop up with a ready to restart button will appear. The inclusion criteria for the second session are 1. No more than 3 resets due to inactivity and 2. Either at least 500 points in at least 3 out of the 20 games, or that the average of games 17 to 20 is at least 100 points higher than the average of game 1 to 4. These criteria were put in place to maximize recruiting only players who were sufficiently attentive and showed signs of learning. Using those criteria, 66 number of players who finished the first session were excluded from participating in the second session. Recruitment continued until 20 participants per condition successfully completed both sessions.

Behavioral results

Figure 3 displays the points per game for each condition; the following analyses will focus only on data from the human players (in red). To get a measure of test performance for each participant, we averaged each participants points across their 8 games of the test block (games beyond the rightmost dashed line in Figure 3).

We then fitted a linear regression with average test points as the dependent variable. The independent variables of interest were training group (consistent or varied) and the thrust type on block 4, which is the last practiced thrust type before the transfer test (high or low). To account for the possible effects of video gaming experience and other participant characteristics on transfer performance, we included the following as nuisance variables: age, gender, the dominant hand used to control movement in games, and the number of hours per week spent on different genres of video games.

Variable training outperforms consistent training on transfer test

From the results of the regression (adjusted $r^2 = 0.32$), varied training ($\beta = 265.46$, Std. Error = 103.11, $p < 0.05$), low thrust on block 4 ($\beta = -272.38$, Std. Error = 98.08, $p < 0.05$), hours per week spent on 2D action ($\beta = 204.88$, Std. Error =51.78, $p < 0.05$) and 3D shooter games ($\beta = 111.04$, Std. Error = 37.02, $p < 0.05$) significantly predicted test points. Notably, players who received varied training were predicted to outperform their consistent counterparts on the transfer test by 265.46 points. This advantage of varied training is aligned with the variability of practice hypothesis.

Changes in performance depends on the direction and magnitude of switch in thrust level

Thrust type on block 4 also strongly predicted transfer performance ($\beta = 185.521$, Std. Error = 76.376, $p < 0.05$), where participants trained on high thrust outperform those trained on low thrust. While we did not predict an effect of recent thrust level, it might be that training with higher thrusts is more difficult and that switching to lower thrust levels is akin to switching to an easier task, which has been shown to facilitate transfer (e.g. Barch and Lewis, 1954).

If there were behavioral differences between games of different thrust levels, one would expect the largest differences to manifest when initially switching to a new thrust. Hence, to further investigate the effect of switching thrust levels, we analyzed the point difference obtained by subtracting the points earned on the last game of a block from the points earned on the first game of the subsequent block. Point differences are then sorted by switch type. For instance, the point difference between games 33 and 32 for a HHHHM participant would be considered a H to M switch, whereas the point difference between the same numbered games for a LLLLM participant would be considered a L to M switch. Point differences for H to L (games 8 to 9 and 24 to 25 for HLHLM and games 16 to 17 for LHLHM) and L to H (games 16 to 17 for HLHLM and games 8 to 9 and 24 to 25 for LHLHM) were gathered from participants in both varied conditions. Switch types were then re-coded as thrust differences to express a quantitative difference in thrust levels (L to H = 0.6 0.2 = 0.4; L to M = 0.2; H to M = -0.2; H to L = -0.4).

A regression model (adjusted $r^2 = 0.3308$) with thrust difference as the sole predictor of point difference estimated a slope of -949.45 (Std. Error = 95.02, $p < 0.05$). This suggests that increasing thrust by 0.2 would result in a drop of 189.89 in points, providing further evidence that switching from low to high thrust decreases performance while switching from high to low thrust increases performance.

Adaptive Control of Thought – Rational

In a recent study, Anderson et al. (in press) demonstrated that an ACT-R model produced the same learning trajectory as humans do (r = 0.96) in a Space Track task where players would play 40 games at a thrust level of 0.3. There are four key features of ACT-R that enabled this successful simulation. First, there are limits on various human cognitive processes such as attention and response times that constrain how human skill acquisition proceeds. The ACT-R cognitive architecture incorporates realistic performance constraints on the speed and accuracy of perception and action. Second, human participants do not begin learning from scratch, but are informed by explicit instructions about the controls and goals of the task. Through instruction following, ACT-R models also utilize task knowledge to accelerate learning in the initial stages. Third, the improvement in performance with experience is partially governed by increased automaticity and faster deployment of knowledge. ACT-R models capture this by production compilation, a process that gradually proceduralizes declarative knowledge and reduces the time cost of having to retrieve declarative knowledge for action execution. Fourth, human skill mastery also relies on tuning the control parameters of ones actions to predictors of success or failure in the task environment. This is captured in ACT-R by a new Controller module that explores continuous dimensions of performance to identify how to control actions. For instance, one dimension that was explored in Space Track was the speed of the ship that would yield an optimal trade-off between number of segments cleared versus ship crashes.

An open question about the new controller module is whether it responds to environmental changes in the same way humans do. Thus, it becomes of interest to see how it responds to the changes investigated in our experiment.

Control Tuning

Through practice, the model learns the optimal values for 5 control variables: aim, ship speed, thrust duration, when to start making a turn, and the ship's orientation with respect to the angle of the upcoming intersection. For each control variable, the model samples values within a preset range and evaluates the mean rate of return for the sampled values according to relevant feedback. Using that feedback, the module then estimates a quadratic function that describes the relationship between rate of return and control value, which in turn influences how the module samples the next set of control values to try. This process is repeated iteratively throughout the experiment, and the model eventually converges to a truer estimate of the relationship between return and control values.

For the model, relevant feedback comes from two sources: the number of crashes and the number of segments cleared. The weights on these sources determine the contribution of each source of feedback to the estimated rate of return. Different source weights potentially relate to differences in risk attitudes; for instance, a player might adopt a riskier approach, clearing more track segments but also crashing more often than a more cautious player.

For our first set of model simulations, we compared models with different weight ratios on the control variables. The reference model weights both features equally (-1 for a crash, +1 for a cleared segment). One modified model reflects the difference in point values assigned by the game to these features and weights a crash four times the benefit of clearing a segment (-4 to +1). Another modified possibility reflects a loss aversive player by weighting crashes as being eight times a cleared segment (-8 to +1).

While points are the primary indicator of performance on the task, two players could conceivably achieve the same total points through different strategies. For instance, a player might adopt a riskier approach to the game, clearing more track segments but also crashing more often than a more cautious player. To further investigate how switching thrust types influences more fine-grained behavior in both models and humans, we also analyzed the mean speed, number of segments cleared, and the number of crashes per game. For each point of comparison, we obtained the sum of squared errors (SSE), which measures the absolute deviation the average model exhibits with respect to the average human across all 4 conditions and 40 games. These results are presented in Table 1.

Different ratios of good and bad weights

The first set of comparisons comprise of the following models: the base (reference) model with a weight ratio of 1:1, a model with a ratio of 4:1, and a (loss aversive) model with a ratio of 8:1. Of the three models, the worst performing model by far on all measures is the base model. The other two models perform comparably, with the loss aversive model performing slightly better than the 4:1 model on all measures except total points earned. The relatively small differences in model fits possibly suggest that the weighting function of some human players might be best characterized by the 4:1 ratio, which reflects the corresponding contributions of crashes and segments cleared to the total points earned, while the weighting function for other players might be better characterized a the 8:1 ratio, which reflects a disproportionately heightened sensitivity to crashes over segments cleared. For the sake of simplicity, we proceeded to incorporate the 4:1 ratio in our subsequent model simulations.

Figure 3 displays how the average points change as a function of game number for both humans and the 4:1 model in all four training conditions; notice that the model shows the same increases and decreases in performance when the thrust level switches as do humans.

Adding slowdown and a decay on past experiences

While the 4:1 ratio model does qualitatively simulate human behavior adequately on the number of crashes, segments cleared, and the overall points earned (refer to Figures 3 and 5), it does a poorer job of capturing how human players modulate their mean speed across games. Referring to Figure 6, it appears that the model drastically changes its speed whereas human players only make small changes in response to changes in the thrust level. This then motivated the next set of models, where we added slowdown, the ability for the model to actively reduce the spaceship's speed when it overshoots its desired control speed value.

Another model manipulation we investigated was to have the model discount its old experiences. Because the Space Track task used in Anderson et al. did not change its parameters over time, it was unsurprising that a model that weighted all experiences equally would be able to perform comparably with those that discounted old experiences. However, as the task used in the current study does introduce changes in the task parameters, it might be reasonable to expect that a model that decays the weight of old experiences would be able to adapt better to the changes in thrust level. When the task parameters change, it is likely that the same control value will result in different payoffs. For instance, pressing the thrust key for 1 second in a low thrust level will increase the ship's velocity by a smaller amount than pressing the key for 1 second in a high thrust level.

There is evidence from the memory literature for an exponential decay function on the retention of past items in memory (e.g. Rubin et al., 1999). Hence, we chose to discount the weight of a past experience by $.995^t$ where t is the time in seconds.

The second set of comparisons comprise of four models: the 4:1 weight ratio model, which also serves as the reference model for this comparison, an exponential decay model, a model with slowdown, and a model with both slowdown and an exponential decay.

Between the four models, there are two that best fit the human player data; the exponential decay model for total points ($SSE = 5594241$) and crashes ($SSE = 1064$) , and the slowdown and decay model for segments cleared ($SSE = 2564$) and mean speed ($SSE = 15.2$). The presence of a decay function in both models suggests that human players might adapt to their current thrust level by discounting the weight of their past experiences, especially if those experiences were obtained from a different thrust level.

Of the four measures, a model's match on total points is the least important because the total points earned is a composite of the segments cleared and crashes measures. Focusing on the other three measures, the slowdown and decay model appears to be the overall better model, especially because the pure decay model's advantage over the slowdown and decay model in crashes appears to be relatively smaller than its disadvantage in segments cleared and mean speed.

Referring to Figures 4 and 6, the slowdown and decay model shows a large reduction in both the number of segments cleared and its mean speed compared to the reference (4:1) model. This reduction is particularly apparent during the high thrust level blocks, and enables the slowdown and decay model to align better with the mean speed of human players across all conditions. Despite the model's success, it should be noted that the model still exhibits larger modulations in its mean speed in response to changes in thrust level than human players do, suggesting that human players might actively aim to maintain the ship speed within a range of values instead of completely adapting the ship speed to optimize the points earned in games of different thrust levels.

Table 1: Model comparisons

Model	Weight Ratio	Slowdown	Decay	*Sum of Squared Errors (SSE)*			
				Total Points	Segments Cleared	Crashes	Mean Speed
Base	-1 : +1			25647698	11064	5087	119.8
Weight=4	-4 : +1			6789958	9413	1370	69.7
Weight=8	-8 : +1			7398453	7383	1299	54.2
Slowdown	-4 : +1	✓		19404654	3178	1516	15.5
Decay	-4 : +1		Exponential	5594241	11660	1064	82.1
Slowdown + Decay	-4 : +1	✓	Exponential	15901165	2564	1327	15.2

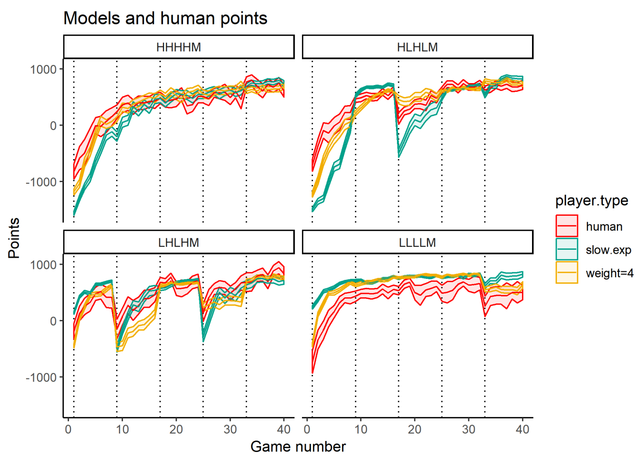

Figure 3: Points by game number for each training condition. Mean human points are in red (n=20 per condition); mean slowdown + decay model points are in green (n=100); mean -4 : +1 weight ratio model points are in yellow (n=100). Shaded areas are S.E.M. Dashed lines indicate the start of a new block.

Figure 4: Segments cleared per game.

Conclusion and Further Work

The behavioral results suggest that the variability of practice hypothesis extends beyond simple motor skills to more dynamic and complex skills that require integrating perceptual, motor and cognitive components. However, our results also indicate that a person's performance on a new thrust level is influenced by their most recently experienced thrust level. Thus, transfer performance depends not only on whether one receives consistent or varied practice, but also on the specific parameters within a consistent or varied training schedule.

Switching from a high thrust to a low or medium thrust improves performance while switching from a low thrust to a high or medium thrust decreases performance. However, it is unclear why these switch effects occur, and why there is an asymmetry in these effects. One possible extension involves investigating how different thrust levels affect motor variation. Motor learning often involves the minimization of motor variation such that one is better able to precisely execute an intended action (refer to Dhawale et al., 2017 for a review). In a high thrust game, a small deviation in the duration of a thrust key press from the intended duration would cause the ship to slow down or speed up more drastically than for the same deviation in a low thrust game. Hence, it might be that players trained on high thrust games have more pressure to control and minimize their motor variation. When switching to a lower thrust level, these players easily adapt to the new thrust because they can immediately apply a suitable degree of control on their thrust key presses. In comparison, players trained on low thrust games have less pressure to minimize motor variation. When these players switch to a higher thrust level, they would be forced to grapple with learning a level of control that was previously unnecessary.

Our model comparisons reveal that the best fitting ACT-R models weight negative events more severely than positive ones. As players are rewarded depending on how many points

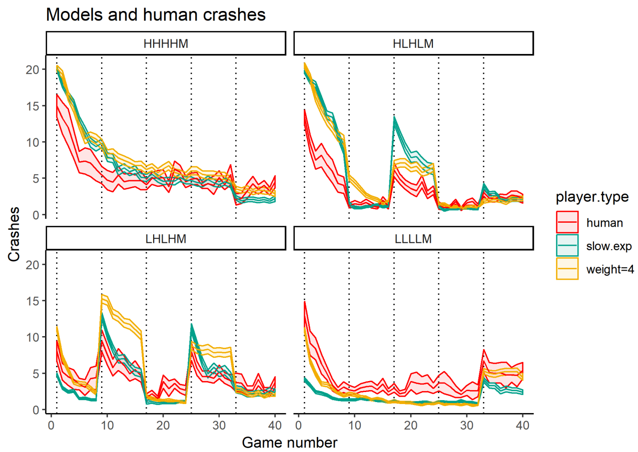

Figure 5: Crashes per game.

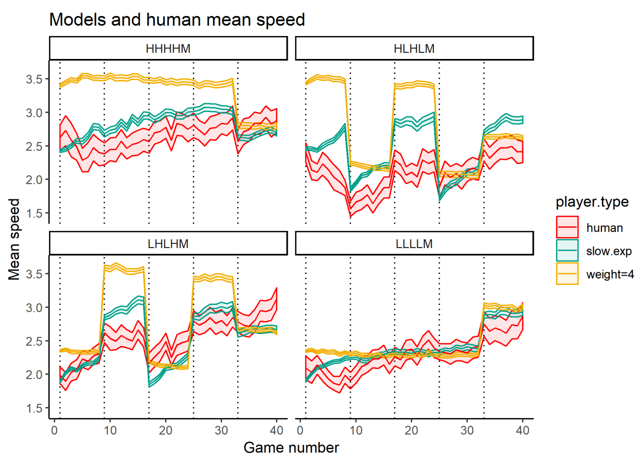

Figure 6: Mean ship speed per game.

they earn, it is reasonable that some players would weight avoiding crashes over clearing track segments in a ratio that reflects their relative contribution to points. Alternatively, as humans have been shown to demonstrate loss aversion in the face of equally valued gambles (e.g. Kahneman and Tversky, 1979), it is also reasonable that some players would place an even greater emphasis on avoiding crashes. Future work would involve investigating if the variability between individual players could be explained by models with different weight ratios.

Our comparisons also provide evidence for including a decay on past experiences. As different thrust levels likely result in different payoffs for the same control setting, a player that discounts old experiences from a previous thrust level would update their estimated payoffs faster when adapting to a new thrust level. More generally, adaptation to changes in the environment is facilitated by prioritizing information learned from recent experiences as these would better reflect the state and reward structure of the current environment.

Finally, while adding the ability to slowdown does improve the models' fit to the human players' mean ship speed, the models still exhibit larger modulations than humans players do when switching between thrust levels. One possibility is that human players are not using points as feedback for ship speed but perhaps using some sense of a comfortable speed. Further work needs to be done to see whether maintaining desired speed can be used as a feedback for the Controller module. Speed control was successfully used as a feedback signal in another video game, YouTurn, described in Anderson et al.; that YouTurn model used speed control to tune one control variable, while using point-related measures to tune other control variables.

Acknowledgments

This research was supported by ONR Grant N00014-15-1-2151 and AFOSR/AFRL award FA9550-18-1-0251.

References

Anderson, J. R., Betts, S., Bothell, D., Hope, R., & Lebiere, C. (in press). Learning rapid and precise skills. *Psychological Review*.

Braun, D. A., Aertsen, A., Wolpert, D. M., & Mehring, C. (2009). Motor Task Variation Induces Structural Learning. *Current Biology*, *19*(4), 352–357.

Dhawale, A. K., Smith, M. A., & Ölveczky, B. P. (2017). The Role of Variability in Motor Learning. *Annual Review of Neuroscience*, *40*(1), 479–498.

Goodwin, J. E., Eckerson, J. M., Grimes, C. R., & Gordon, P. M. (1998). Effect of Different Quantities of Variable Practice on Acquisition, Retention, and Transfer of An Applied Motor Skill. *Perceptual and Motor Skills*, *87*(1), 147–151.

Kahneman, T. A., Daniel. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263–292.

Kerr, R., & Booth, B. (1978). Specific and varied practice of motor skill.pdf. *Perceptual and Motor Skills*, 395–401.

McGovern, D. P., Webb, B. S., & Peirce, J. W. (2012). Transfer of perceptual learning between different visual tasks. *Journal of Vision*, *12*(11), 4–4.

Rubin, H. S. C., David. C., & Wenzel, A. (1999). The precise time course of retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(5), 1161–1176.

Speelman, C. P., & Kirsner, K. (2001). Predicting transfer from training performance. *Acta Psychologica*, *108*(3), 247–281.