Establishing a paradigm to investigate strategy use in complex skills

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

Keywords: strategy use; skill acquisition; practice; decision making; adaptivity; ACT-R

Introduction

Questions of strategy selection have been studied in various contexts such as problem solving, text editing, and even dynamic, fast-paced tasks. One way to model the strategy selection process is as a learning and decision problem: with experience, the agent learns the expected utilities of strategies, and executes a strategy based on what it has learned (Lovett & Anderson, 1996).

It is important to note that the strategies studied in most of the past research have relatively stable utilities. Even when the task structure is manipulated to change the utilities of strategies, these changes are relatively infrequent (Schunn & Reder, 2001). This contrasts with many real-world skills, such as sports and video gaming, where different strategies are optimal at different points during the learner's trajectory. As a learner practices a skill, improvements in the learner's degree of perceptual-motor calibration to the physics of tools and devices interacts with the difficulty of executing a strategy to affect the strategy's utility. Furthermore, it is often unknown what the maximum utility of any strategy will be, as this is partly determined by the learner's own general perceptual-motor abilities and prior experiences.

How humans learn and select strategies in the face of such variation and uncertainty behooves further investigation. Towards that goal, we present a task and strategy paradigm that captures many of the features of a typical complex skill. We also examine possible interactions between strategy use, perceptual-motor calibration, and task knowledge using past experimental data and model simulations within the Adaptive Control of Thought-Rational (ACT-R) cognitive architecture.

Space Track

Space Track is a video game developed by Anderson et al. (2019). The player's goal is to earn points in 3-minute trials by navigating a spaceship along a racetrack comprising of rectangular segments (Fig. 1). Completing a rectangle awards 25 points, while crashing into the track walls loses 100 points. Players control the ship using three keys: "A" and "D" to rotate counter-clockwise and clockwise, and "W" to thrust.

Space Track is similar to many sports and video games in three ways. It simultaneously engages visual, motor, and cognitive processing, requires the rapid and precise execution of actions, and has a relatively high performance ceiling. In the study by Anderson et al. (2019), the average human scores around 900 points by the end of 40 trials. This is far under 2350, the highest number of points achieved by one subject.

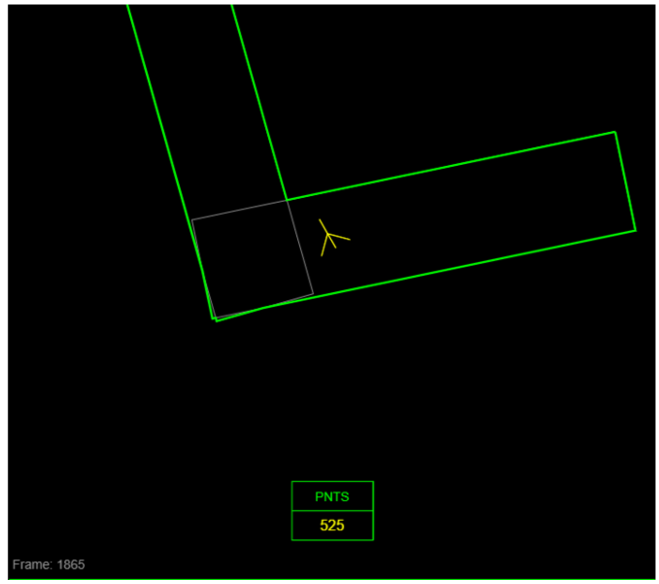

Figure 1: An example Space Track screen

Turning and Calibration

Successful navigation relies on learning the physics of the game. Unlike real-life driving, the racetrack in Space Track is frictionless. To get a car to travel along some desired trajectory, a driver would orient the front wheels to align with that trajectory and then accelerate. If a naïve turner attempts the same in a frictionless space, they will instead send the ship careening away from the desired trajectory due to the residual velocity from the previous trajectory. Turning in frictionless space requires turning less than the desired angle and thrusting until the ship is flying in the desired direction. For any desired speed along the new trajectory there is a unique angle of under-turn that achieves it. Skilled players seem to have learned this angle and how long to thrust.

Through practice, players also gain experience calibrating their perceptual-motor system to account for the continuous and dynamic nature of the game. Human perceptual-motor processing requires time, with an average minimum of 190ms between detecting a visual scene and executing a keypress (Woods et al., 2015). Even a delay of 100 ms between detecting the ship's orientation and lifting the finger from the turn key results in an additional 18 degrees of rotation. Thus, the player needs to account for that lag time by learning the appropriate visual cues for beginning an action. Similarly, a player needs to learn how close to the desired new angle they should be before lifting their thrust finger.

Stopping as a strategy

One potential strategy we identified from prior experiments (Seow et al., 2019) involves rotating the ship in the direction opposite to its current trajectory and thrusting until the ship comes to a halt. Stopping at track corners compensates for naïve turning, since the resultant trajectories from naïve and optimal turners are identical when the ship is stationary. Stopping also gives inexperienced and less calibrated players more time to react to changes in track curvature.

On the flip side, stopping limits average ship speed, which in turn limits the maximum possible distance a player can cover within each 3-minute trial. Thus, it is not immediately obvious whether stopping has a positive or negative utility.

Using data from Seow et al. (2019), we found that increased stopping did not predict a change in the average points earned per trial (Fig 2). Rather, increased stopping correlated with a narrower range of points, raising the floor while lowering the ceiling. This suggests that when scoring is below the mean, increased stopping use might improve performance, but when it is above that mean, increased stopping use might instead lead to worse performance.

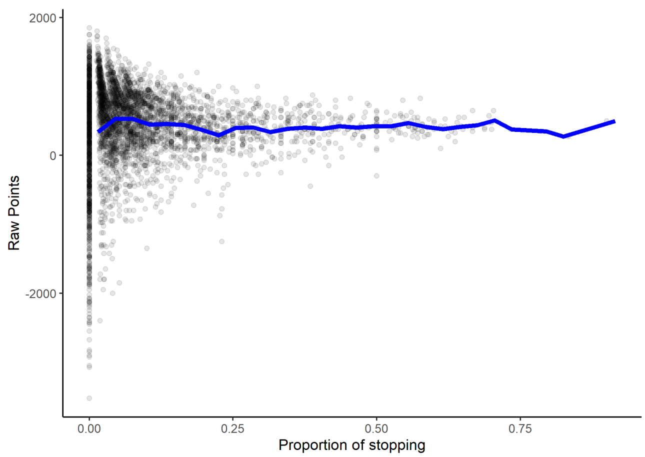

Figure 2: Points per trial in humans. The blue line tracks the average points earned across proportions of stopping.

Stopping, Turn Optimality, and Experience

To test these potential tradeoffs of stopping, we simulated learning on Space Track using ACT-R models. We modeled variations in turn optimality as 11 weighted combinations of the contributions from the naïve and optimal turning algorithms. Differences in stopping use was captured by two types of models, one that stopped at every corner, and one that relied on its turn algorithm to navigate corners. Stopping use was crossed with turn optimality to yield 22 models.

Stopping helps inexperienced agents and naïve turners but limits experienced agents and optimal turners (Fig. 3). A regression model ($r^2 = 0.74$) further showed that score was predicted by the interaction between stopping, turn optimality, and trial number ($\beta = 2.19$, $SE = 0.65$, $p < 0.05$).

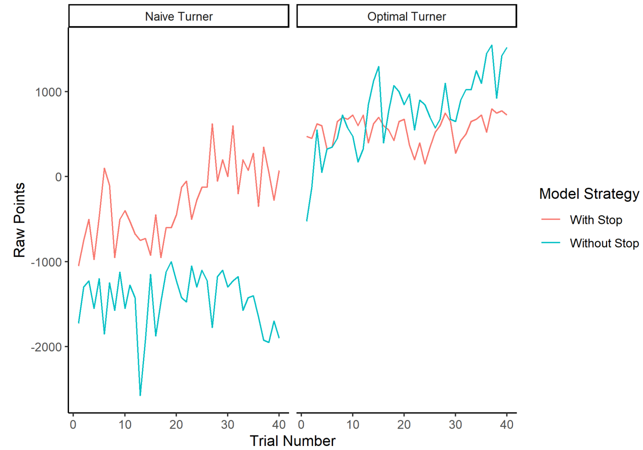

Figure 3: Model points across trials. The stopping model (in red) performs better than the non-stopper except when the agent is an optimal turner and has had sufficient experience.

Conclusion and Further Research

We have identified a task and strategy paradigm that is potentially suitable for understanding the processes that underlie the learning and selection of strategies in complex skills. With Space Track and the strategy of stopping, one promising future direction is to apply and test current models of decision making (e.g. Reinforcement Learning) on human gameplay to investigate how basic decision processes account for strategy shifts in complex skill acquisition.

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. (2019). Learning rapid and precise skills. *Psychological review*, *126*(5), 727.

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive psychology*, *31*(2), 168–217.

Schunn, C. D., & Reder, L. M. (2001). Another source of individual differences: Strategy adaptivity to changing rates of success. *Journal of experimental psychology: General*, *130*(1), 59.

Seow, R., Betts, S., & Anderson, J. R. (2019). Transfer effects of varied practice and adaptation to changes in complex skill acquisition. In *Proceedings of the 17th international conference on cognitive modelling* (p. 222-227).

Woods, D. L., Wyma, J. M., Yund, E. W., Herron, T. J., & Reed, B. (2015). Factors influencing the latency of simple reaction time. *Frontiers in human neuroscience*, *9*, 131.