Cognitive Systems Research 86 (2024) 101239

Contents lists available at ScienceDirect

Cognitive Systems Research

journal homepage: www.elsevier.com/locate/cogsys

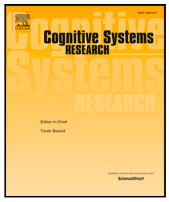

Tight resource-rational analysis☆

Cvetomir M. Dimov [a,b,*], John R. Anderson [a], Shawn A. Betts [a]

[a] Department of Psychology, Carnegie Mellon University, United States of America
[b] Department of Psychology, University of Geneva, Switzerland

ARTICLE INFO

Action editor: S Thill

Keywords:
Rational analysis
Tight resource-rational analysis
Dynamic task
Space track

ABSTRACT

Resource-rational analysis is used to develop models that assume that people behave optimally given the structure of the task environment and the cost of cognitive operations. We argue in favor of a tight resource-rational analysis, an extension in which model parameters are independently constrained. As a case in point, we demonstrate how to develop a tight resource-rational model of the video game Space Track. Our approach consists of four steps. First, we measure performance-critical parameters in independent micro-tasks, which we input into mathematical models of cognitive processes. Second, we validate these models in other process-specific micro-tasks. Third, we rely on a theory of the cognitive architecture (i.e., ACT-R) to derive estimates of the time costs of these processes. Finally, we generate predictions for the main task, Space Track, by assuming that subjects are doing their best given their abilities. The generated individualized predictions were close to observed subject asymptotic performance, which demonstrated the viability of our approach, even in tasks of similar complexity to that of Space Track.

1. Introduction

Ever since rational analysis was first introduced (Anderson, 1991; Anderson & Milson, 1989), evidence has been accumulating that humans are adaptively rational, that is, optimally adapted to their surroundings. People have been demonstrated to achieve near-optimal performance in perception, motor control, inductive and statistical learning, and reasoning in their natural environment (see, Lieder & Griffiths, 2020, for a review). For some tasks, a careful task analysis sufficed to demonstrate optimality, without considering cognitive constraints. For example, Oaksford and Chater (1994) demonstrated that, although disagreeing with the normative prescription, people's choices in the Wason selection task maximized information gain in an environment where causal relations are rare. Optimality without considering cognitive limitations was also showcased for sequential decision making in uncertain environments (Sims, Neth, Jacobs, & Gray, 2013) and even for curiosity (Dubey & Griffiths, 2020), among others.

Yet, when applying rational analysis to other tasks, details about the cognitive architecture were crucial in predicting task performance and these were named *resource-rational* (Lieder & Griffiths, 2020). Resource-rational models are typically more complex, because they include a characterization of both task and cognitive capacities. For example, let us take Howes, Lewis, and Vera (2009)'s *cognitively bounded rational analysis* of the psychological refractory period. The psychological refractory period is a slow down of responses when two choice reaction time tasks are in close temporal proximity. Howes et al. (2009) investigated if the slow down is a consequence of architectural or strategic limitations. To address this question, they incorporated two theories of the mind (the cognitive architectures ACT-R, Anderson, 2007, and EPIC, Meyer & Kieras, 1997) into their analysis and predicted the refractory period for multiple strategies. Hence, even a task that appeared simple at first sight resulted in a relatively complex model.

Complex models contain multiple parameters with varying parameter ranges input into functions of different form (Pitt, Myung, & Zhang, 2002). Complex models are naturally capable of fitting a broad spectrum of data patterns, which raises concerns about such models overfitting the data. Confidence in such models will depend on getting the model details sufficiently correct or, as Lieder and Griffiths (2020) concretely put it, resource-rational models need to be optimal "with respect to [people's] actual cognitive constraints" (p. 13). To ensure that the underlying constraints are correct, they suggest that "independently measuring people's cognitive constraints [is] an important direction for future work" (same page).

Here we demonstrate how to achieve this. Our approach marries resource-rational analysis with a proposal put forth half a century ago by Newell (1973) in one of his three suggestions on how to build stronger psychological theories. Specifically, he proposed to "focus a

☆ This work is supported by AFOSR/AFRL award FA9550-18-1-0251.
* Correspondence to: Campus Biotech, Chemin des Mines 9, CH-1202, Geneva, Switzerland.
E-mail address: cvetomir.dimov@unige.ch (C.M. Dimov).

https://doi.org/10.1016/j.cogsys.2024.101239
Received 12 September 2023; Received in revised form 31 March 2024; Accepted 5 April 2024
Available online 12 April 2024
1389-0417/

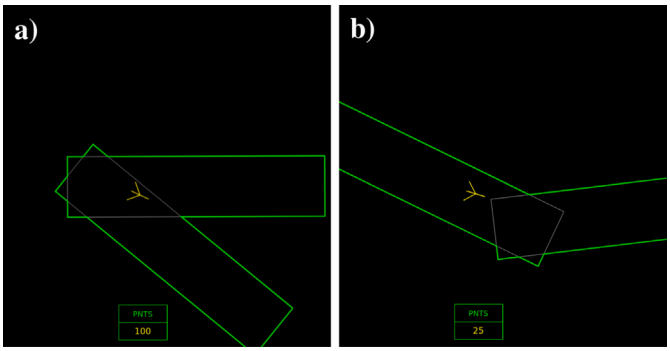

Fig. 1. The dynamic task Space Track. The task goal is to traverse as many rectangular segments (25 points), while crashing as little as possible (−100 points). Segments are oriented at various angles relative to each other, randomly drawn from a uniform distribution between 30 (panel a) and 150 degrees (panel b).

series of experimental and theoretical studies around a single complex task" with "all the studies… designed to fit together…" (p. 303). Newell gave an example of a doctoral dissertation on mental multiplication, in which all model parameters were estimated in separate micro-tasks, and the thus constrained model tested on the main task.

In this work, we will strive to adhere to Newell's prescription when constructing a resource-rational analysis of a complex task. Specifically, we will use a series of independent micro-tasks to estimate per-subject model parameters and to validate model components. We will then input the estimated parameters and validated components into a resource-rational model of a main task. The model will incorporate task-relevant components of a theory of the mind (i.e., ACT-R), the measured individualized parameters, and the assumption that each subject is doing their best. This will result in an individualized no-free-parameter resource-rational model, what we call, to use Newell's term, a *tight* resource-rational model. We will now describe the main task that we focus on in this work before proceeding with the experimental design and model development.

1.1. The main task

The task of choice, Space Track, is a dynamic task recently used in studies and modeling endeavors on skill acquisition (Anderson, Betts, Bothell, Hope, & Lebiere, 2019; Seow, Betts, & Anderson, 2021). In this task, subjects fly a spaceship through a frictionless track in space that is constructed of equally-sized rectangular segments, oriented at various angles β relative to each other (uniformly distributed between 30 degrees, Fig. 1, a, and 150 degrees, Fig. 1, b). As soon as the ship traverses one segment and enters a new one, the subsequent segment appears. The goal is to fly through as many segments as one can while crashing as little as possible, as each traversed segment brings 25 bonus points and each crash leads to a 100 point penalty. The ship can turn clockwise ("D" key) and counterclockwise ("A" key), and thrust forward ("W" key). The effect of pressing the keys is proportional to the length of the key press. For example, pressing the thrust key longer will result in a larger change in velocity.

A ship in a frictionless environment behaves differently than a car on the ground. For example, the ship maintains its velocity unless a thrust changes it. Also, the ship does not fly in the direction that it is facing. Instead, its velocity after a thrust is a vector sum of its starting velocity and the velocity change introduced by the thrust. Hence, to steer the ship, a subject needs to take the ship's current velocity into account in order to aim the ship appropriately and then thrust. Learning how and when to do this is one of the difficulties that subjects face in this task (those interested, can try the task out at https://andersonlab.net/demos/small-tasks-spacetrack/).

Although initially counterintuitive, subjects eventually learn to control the ship. The ACT-R model of the task (Anderson et al., 2019) describes how this learning process takes place, which involves learning multiple control parameters and proceduralization of the knowledge of how to act in the task. Asymptotically the model reaches the procedural phase of skill acquisition, in which it directly maps the task state to key presses.

Central for our experimental design is that, according to this model, the key presses can be executed in two ways: a single short ballistic motor act that we call a *tap* or two consecutive motor acts (i.e., a key press act followed by another motor act - release of that key) that together constitute a *long key press*. We expect that taps and long key presses have different properties, because tap accuracy should only be limited by the motor system, whereas long key press accuracy by both perceptual and motor systems. For example, during a long turn, subjects decide to release the pressed key when they detect that the ship is at a certain orientation (perceptual system), which is then followed by the actual key release (motor system).

These key presses are usually executed towards a certain goal such as a target ship orientation or a target ship velocity. Reaching such a goal might involve one or several key presses (either taps or long key presses). We call a sequence of key presses toward a single goal an *action*. Since both taps and long key presses are involved in actions, individual differences in both types of key presses should determine turn or thrust action properties.

Although paralleling each other in many ways, turn and thrust actions differ in one key manner: turning is bidirectional, whereas thrusting is not. This means that turn errors can be immediately corrected with further turn key presses, whereas thrust errors can only be corrected if the thrust had been shorter than intended. If longer, a thrust can be corrected only if the ship is re-aimed before thrusting again, that is, an additional turn action is needed. This means that turn actions have the potential to be more accurate, but also that they might require more key presses if subjects execute many corrections on the way to achieving that accuracy.

Action properties, such as their accuracy or the number of key presses required to complete the action, influence Space Track performance. For example, noise in target ship orientation and in target velocity, which is inevitably introduced by the actions, will offset the ship from its ideal trajectory. This, in turn, will affect the probability that the ship crashes. Additionally, the more slowly subjects complete their actions (i.e., with more key presses), the less likely they are to steer the ship away from an impending crash. Note that action completion time includes both the time to execute each key press and any cognitive delays associated with preparing the key presses. The cognitive delays include not only motor preparation time, but also those related to perceptual and central cognitive processes.

To summarize, mastering Space Track consists in mastering two actions, the accuracy and latency of which are a function of multiple cognitive processes. This makes the task a good candidate for showcasing our approach. Moreover, previous studies with this task have demonstrated large individual differences in performance, which suggest a large explanatory potential based on individual difference measures. A rational analysis of this task would reveal if Space Track performance is determined mostly by subjects' individual perceptual-motor limitations or if the learning process introduces suboptimalities in performance (as has been demonstrated in other complex tasks, see Fu & Gray, 2006), because if subjects were to reach the performance predicted by an optimal model, that would indicate that the learning processes have impaired performance very little.

2. Methods

2.1. Participants

Subjects were recruited on Amazon Turk. Informed consent approved by the Carnegie Mellon University Institutional Review Board

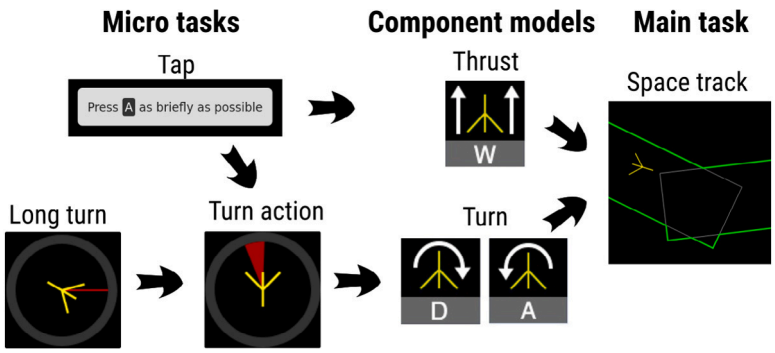

Fig. 2. Study design. Two micro-tasks (Tap task and Long turn task) serve to measure tap duration and long key press accuracy (specifically of turning). These are used to develop models of the two actions in the main task: turning and thrusting. The turning model is validated in another micro-task (Turn action task). The turn and thrust action models are the building blocks of a rational analysis of the main task, Space Track.

was obtained from each participant. 18 subjects (13 males, mean age: 33 years, min age: 25, max age: 43) participated in this study.[1] Three subjects were excluded from further analyses because of their poor performance in one of the tasks (2 in the Tap task and 1 in Space Track, see below). All subjects were paid a base payment of $20 and a performance-contingent bonus.

2.2. Procedure

The study consisted of a series of micro-tasks, followed by the main task, Space Track (Fig. 2). Two micro-tasks served to measure tap (*Tap task*) and long key press accuracy (specifically of turning, *Long turn task*). We only included a long key press task for turning, because we expected that they have a larger influence on Space Track performance than thrust actions. A third micro-task served to develop and validate a model of turn actions (*Turn action task*), which was necessary because we could not validate the turn action model directly in the main task.

We expected that the parameters measured in the micro-tasks would be relatively stable over the time course of the experiment (~2 h) and, hence, be accurate estimates of perceptual-motor constraints in Space Track. Note that all tasks include rewards to incentivize best performance. Reward in the main task is linear with performance, where segment traversals gain a quarter of what crashes cost (25 vs. −100 points). Reward in all micro-tasks is a quadratic function of performance: the better subject performance, the superlinearly more points they receive (see Appendix A for details). We now proceed by describing all tasks.

2.2.1. Tap task

The Tap task was designed to measure the mean $\bar{t}_{\text{tap}}$ and variability σ_{tap} of tap duration. At the beginning of each trial, subjects were directed to press one of three keys ("A", "D", and "W") for as short a duration as they could (see Fig. 2). Immediately after the key press was completed, subjects received feedback about the duration of the key press (in ms) and the corresponding trial reward. The task consisted of 150 trials and took about 13 min to complete on average.

[1] We expected a correlation between the micro-task measures and Space Track of at least 0.77 based on inter-task correlations in Anderson et al. (2019). Specifically, in that study, subjects were faced with the tasks Space Track and You Turn, each of which was played for 20 games. The correlation between performance in the two tasks was 0.77 (coincidentally, both in the "Space Track - You Turn" condition and the "You Turn - Space Track" condition). To detect such a correlation with $\alpha = 0.05$ and $\beta = 0.20$, a sample size of 11 is needed. Conservatively, we collected data from a few more subjects.

2.2.2. Long turn task

The Long turn task was designed to measure the variability of a long turn key press σ_{tu}. On each trial, the ship started at a random orientation α and its target orientation α_{tar} was denoted by a red line. Subjects could press a single key (either "A" or "D") only once to orient the ship as closely as possible to the target orientation. The ship's starting orientation was chosen among 12 values evenly spaced between 0 and 330 degrees (i.e., every 30 degrees). Starting orientations were paired with 10 evenly spaced values between 90 and 180 degrees of relative target orientation, which was randomly selected to be either clockwise or counterclockwise relative to the starting orientation. Large relative target orientation values were chosen so that turns cannot be completed in a single ballistic key press (e.g., at the given turn rate, turning 90 degrees takes 500 ms, which is roughly an order of magnitude longer that the duration of taps that subjects are capable of). The exhaustive pairing of starting and relative target orientations resulted in 120 trials. Moreover, an additional 10 practice trials were added at the beginning for a total of 130 trials. Subjects took about 15 min on average to complete them.

2.2.3. Turn action task

In the Turn action task, subjects had to orient the ship within a region ε_α around a target orientation, that we call *aim accuracy*. Unlike in the Long turn task, subjects were free to use as many key presses n_{tu} as they desired, but the trial was only completed once the ship was within the target region. Subjects were rewarded based on the time $t_{\text{tu act}}$ they needed to complete the turn action. Two different aim accuracy ε_α values were used: ±6 degrees and ±12 degrees. As in the Long turn task, the ship's starting orientation was chosen among 12 evenly spaced values between 0 and 330 deg, whereas its relative target orientation was chosen among 6 evenly spaced values between 90 and 170 degrees (i.e., every 20 deg). Additional 10 practice trials were added at the start, for a total of 154 trials. This task took about 15 min, on average, to complete.

2.2.4. Main task: Space track

Subjects started by reading detailed instructions about the game mechanics, actions and goal, by the end of which they had full knowledge of how the game is played. They continued with 20 3-minute long games of Space Track and finished with a feedback questionnaire. Subjects could see their score at any moment during the game. They took, on average, 77 min to complete this task.

3. Results and discussion

We estimated asymptotic performance for each task by removing all learning trials and outliers (see Appendix B for additional details

Table 1
Ranges of individual means and standard deviations for asymptotic values on micro-tasks.

Task	Measure	Notation	Min	Max	Median
Tap	Mean duration	$\bar{t}_{\text{tap}}$	20	105	57
	SD of duration	σ_{tap}	8	51	15
Long turn	SD of error of angle	σ_{tu}	5.4	16.8	7.9
Turn action	Mean number of key presses	$\bar{n}_{\text{tu},\epsilon_\alpha=6}$	1.3	2.7	1.6
	(given tolerance ϵ_α)	$\bar{n}_{\text{tu},\epsilon_\alpha=12}$	1.0	2.3	1.3
	Probability of one key press	$P(n_{\text{tu},\epsilon_\alpha=6}=1)=P_{1,\epsilon_\alpha=6}$	0.30	0.75	0.57
	(given tolerance ϵ_α)	$P(n_{\text{tu},\epsilon_\alpha=12}=1)=P_{1,\epsilon_\alpha=12}$	0.39	0.98	0.74
Space Track	Mean score	Score	−919	1475	494
	Mean number of crashes	n_{crash}	0.25	12	1.75
	Mean speed	$\bar{s}$	0.53	1.79	0.94
	Mean number of key presses in a segment	$\bar{n}_{\text{kp}}$	3.20	8.18	5.14

of the data cleaning procedure).[2] Specifically, we eliminated the first 25 trails for the micro-tasks for most subjects. For two subjects in the Tap task and two subjects in the Long turn task, whose performance was not stable after trial 25, additional data points were eliminated. To remove outliers in the Tap and Long turn tasks, we eliminated all data points outside 1.96 SD around the mean tap duration and turn angle offset, respectively.[3] In the Turn action task, all uncompleted trials were excluded. For Space Track, we computed various performance measures in the four consecutive games that produced the highest average score as the best approximation to asymptotic performance. We selected the best scoring consecutive games instead of the last games, because some subjects systematically decreased their performance toward the end of the experiment, likely because of loss of motivation.

3.1. Individual differences

Subjects demonstrated large individual differences on all measures in all tasks (see Table 1 for a summary of the estimated per-subject quantities; Appendix C presents plots related to each task). In the Tap task, the most capable subject could produce a five times smaller key press duration ($\bar{t}_{tap}$ = 20 ms) than the least capable ($\bar{t}_{\text{tap}}$ = 105 ms). Note that two subjects averaged $\bar{t}_{tap}$ above 200 ms, which likely resulted from misunderstanding the task. As mentioned in the methods section, we excluded these subjects from further analyses, because tap measures are a necessary input of our model (see Appendix H, for a further justification).

This value is effectively the lower limit on the duration of both turn key presses and thrust key presses in the main task. Translated to turn angle (turn rate in Space Track: 3 deg/tick[4]; game tick duration: 16.(6) ms), this corresponded to a minimum angle change as small as $\overline{\Delta\alpha}$ = 4 deg for the best subject and as large as $\overline{\Delta\alpha}$ = 19 deg for the worst. Analogously, translated to speed change (acceleration: a = 0.05 pixels/tick^2), this corresponded to a minimum speed change between $\overline{\Delta s}$ = 0.06 for the best subject and $\overline{\Delta s}$ = 0.32 pixels/tick for the worst.

Large individual differences were observed in the turn-related micro-tasks too. In the Long turn task, there was a threefold difference between the most capable (long turn variability σ_{tu} = 5.4 deg) and least capable subject (σ_{tu} = 16.8 deg). In the Turn action task, we relied on two measures of performance: the average number of key presses needed to complete the trial $\bar{n}_{\text{tu}}$ and the probability of completing the trial in a single key press $P(n_{\text{tu}}=1)$, which we estimated separately for the two tolerance regions (ϵ_α = 6 deg and ϵ_α = 12 deg). The relative differences in performance along those measures were between 2 and 2.5-fold.

We characterized Space Track with mean score, mean number of crashes n_{crash}, mean speed $\bar{s}$, and mean number of key presses in a segment $\bar{n}_{\text{kp}}$. There were large individual differences on all four measures. For example, some subjects could not reach positive score even in the their best 4 games, whereas others scored positively already in game 1.[5] Similarly, there was a more than threefold difference in ship speed and a 2.5-fold difference in the number of key presses in a segment.

3.2. Relationship between measures

Table 2 summarizes the correlations between the measures presented in the previous section.[6] The various micro-task measures correlate highly significantly with each other (lowest r: 0.66, highest r: 0.93). Interestingly, Tap duration and Long turn variability correlate strongly (r = 0.81) even though the first is a purely motor measure, whereas the second perceptual-motor. Additionally, performance measures in the Turn action task correlate with both Tap duration and Long turn variability, since key presses in a turn action can be both taps or long key presses. The high inter-task correlations above point toward a single motor factor behind the performance in all micro-tasks.

Micro-task measures also strongly predicted Space Track score and number of crashes, which corroborates our hypothesis that asymptotic Space Track performance is mostly limited by perceptual-motor skills. However, they did not correlate significantly with average ship speed $\bar{s}$ and average number of key presses in a segment $\bar{n}_{\text{kp}}$. Instead, only Space Track score significantly correlated with $\bar{s}$. Finally, there was a strong negative correlation between $\bar{s}$ and $\bar{n}_{\text{kp}}$. This correlation was likely a consequence of the fact that a faster flying ship traverses a segment in less time, which in turn restricts the number of key presses that can be executed.

4. A tight resource-rational model of Space Track

What is the functional relationship between the parameters measured in micro-tasks and Space Track score? To answer this question, we develop a resource-rational model that starts from tap and long key press parameters as inputs into turn and thrust action models, then propagates the effects of these actions onto a track parameterized as

[2] All data and data analysis code can be found on the Open Science Foundation website at 10.17605/OSF.IO/MGFXH.

[3] Since the outlier removal procedure led to an underestimation of sample variability, SDs were adjusted by dividing them by 0.872. Specifically, when the outlier removal procedure is applied to a normal distribution, the resulting subsample has a standard deviation equal to 0.872 that of the original.

[4] A tick is the time unit of the game. The game runs at a frequency of 60 Hz, which results in a game tick duration of 16.(6) ms.

[5] A ship left to fly freely at the starting speed of 0.5 pixels/tick would crash approximately every 15 s and reach a score of −1200 points. This would be the score that a passive subject would reach. We used this score value as a performance threshold. One subject did not reach even that score in its best consecutive 4 games and was therefore excluded.

[6] Since the shortest key press duration exhibited scalar variability (σ_{tap} was proportional to $\bar{t}_{\text{tap}}$ with a coefficient of 0.35, r = 0.86), we only include $\bar{t}_{\text{tap}}$ in the table.

Table 2
Correlation between individual measures on micro tasks and Space Track. $\bar{t}_{tap}$ - mean tap duration; σ_{tu} - long turn SD; $\bar{n}_{tu,\ \varepsilon_\alpha=x}$ - average number of key presses needed to turn to an aim tolerance of x degrees; $P_{1,\ \varepsilon_\alpha=x}$ - probability of turning to an aim tolerance of x degrees in a single key press; $\bar{s}$ - average speed; $\bar{n}_{kp}$ - average number of key presses in a segment.

	$\bar{t}_{tap}$	σ_α	$\bar{n}_{tu,\varepsilon_\alpha=6}$	$\bar{n}_{tu,\varepsilon_\alpha=12}$	$P_{1,\varepsilon_\alpha=6}$	$P_{1,\varepsilon_\alpha=12}$	Score	n_{crash}	$\bar{s}$
σ_α	.81***								
$\bar{n}_{tu,\varepsilon_\alpha=6}$	.88****	.80***							
$\bar{n}_{tu,\varepsilon_\alpha=12}$	.76**	.79***	.84****						
$P_{1,\varepsilon_\alpha=6}$	−.86****	−.77***	−.86****	−.66**					
$P_{1,\varepsilon_\alpha=12}$	−.87****	−.88****	−.87****	−.93****	.88***				
Score	−.81***	−.91****	−.82***	−.86****	.77***	.88****			
n_{crash}	.85****	.93****	.84****	.84****	−.79***	−.92****	−.91****		
$\bar{s}$	−.48	−.49	−.51	−.58*	.43	.49	.74**	−.39	
$\bar{n}_{kp}$	.02	.05	.05	−.02	−.01	.08	−.30	.40	−.80***

in our experiment, and quantitatively predicts score.[7] The central idea behind this model is that unless subjects have sufficient time at their disposal to complete their actions, they will crash and lose valuable points. So subjects need to adjust their behavior to avoid this while still managing to traverse as many segments as possible. In our model, subjects adjust their behavior through two freely chosen parameters, the target speed they aim to fly at s_{tar} and the aim accuracy ε_α they reach before thrusting. Both of these parameters influence how much time is needed to complete actions, *action completion time* t_{act}, and how much time is left before the ship crashes, *time to crash* t_{crash}. For example a narrow aim accuracy requires more key presses to reach than a wide one, leading to long t_{act}. However, it also leads to a more accurate trajectory and hence longer t_{crash}. Analogously, flying at a high speed needs longer thrusts to steer the ship and, consequently, longer t_{act}, while also leaving short time to crash t_{crash}. However, it also increases the rate at which segments are traversed. The optimal values of ε_α and s_{tar} will be highly individual for each subject.

The Space Track model estimates the optimal score for each subject in 5 steps (see Fig. 3, for an overview). Specifically, the model estimates the t_{act} (steps 1–3) and t_{crash} (step 4) distributions for each subject as a function of the individual parameters measured in the micro-tasks (highlighted in red in that figure), parameters for a typical subject derived from the cognitive architecture ACT-R (highlighted in yellow), and the two free parameters ε_α and s_{tar} (in green). To estimate the action completion time distribution, we first develop models of the number of key presses in turn and thrust actions (step 1). We then generate estimates of the cognitive costs τ associated with these actions derived from models constructed in ACT-R (step 2). By combining these two, we produce individualized distributions of turn completion time $t_{\text{tu act}}$ and thrust completion time $t_{\text{th act}}$ (step 3). To estimate the time to crash distribution, we investigate the affect of aim accuracy ε_α, target speed s_{tar} and thrust variability $\sigma_{\text{th act}}$ on flight direction variability σ_θ and the resultant t_{crash} (step 4). In the final model, we estimate the number of crashes and number of traversed segments by comparing t_{act} and t_{rash} and compute score (step 5). This final model also exploits the symmetry in the problem by assuming periodic boundary conditions (explained below).

We will now provide a brief overview of ACT-R, followed by a detailed explanation of the 5 steps. We will provide examples with three inter-segment angles β throughout, which cover the full range of possible values: $\beta = 30$ deg (flight direction almost reversed), $\beta = 90$ deg (segments perpendicular), and $\beta = 150$ deg (only a small trajectory adjustment necessary). Moreover, in all calculations, we will consider ship target speed between $s_{\text{tar}} = 0.5$ pixels/tick, which is the speed at which the ship starts, and $s_{\text{tar}} = 3$ pixels/tick, which is the maximum speed it can fly at. A overview of all variables used in the analysis can be found in Table 3 for reference.

[7] All model analysis code can be found on the Open Science Foundation website at 10.17605/OSF.IO/MGFXH.

Table 3
Main variables used in the resource-rational analysis of Space Track.

Variable	Notation
Action completion time	t_{act}
Time to crash	t_{crash}
Inter-segment angle	β
Velocity	v
Velocity noise	Δv_{noise}
Speed	s
Speed offset	Δs
Flight direction	θ
Flight direction variability	σ_θ
Aim	α
Aim accuracy	ε_α
Long turn variability	σ_{tu}
Number of key presses in turn action	n_{tu}
Turn completion time	$t_{\text{tu act}}$
Time to prepare first turn	$\tau_{n_{\text{tu}}=1}$
Time to prepare turn correction	$\tau_{n_{\text{tu}}>1}$
Long thrust variability	σ_{th}
Number of key presses in thrust action	n_{th}
Thrust action variability	$\sigma_{\text{th act}}$
Thrust completion time	$t_{\text{th act}}$
Time to prepare first thrust	$\tau_{n_{\text{th}}=1}$
Time to prepare thrust correction	$\tau_{n_{\text{th}}>1}$

4.1. A brief overview of ACT-R

ACT-R is both a unified theory of cognition and a software implementation of that theory, which can be used to construct cognitive models. Here we will provide a brief overview of the cognitive architecture, sufficient to follow the model development below. Those interested in a more detailed introduction to ACT-R can refer to the primer of Dimov, Khader, Marewski, and Pachur (2019) or the detailed book of Anderson (2007).

ACT-R is split into independent modules that model different cognitive capacities. At the heart of ACT-R lies a procedural module, which is responsible for selecting cognitive operations to execute. It stores a collection of production rules (or IF-THEN rules), which fire when the cognitive system matches their conditions (the IF part) and execute the operations specified in their actions (the THEN part). Conditions are states of other modules, whereas actions are requests to those modules. A production rule fires when module states agree with its conditions, which then results in requests to modules to perform certain operations and thus change their state. This can enable other production rules to fire, which can further change module states, and so on.

Modules are responsible for different aspects of peripheral and central cognition. For example, there are perceptual modules responsible for visual and auditory perception. In addition, there are motor modules responsible for interaction with computer devices and for producing speech. Moreover, there is a module that holds the current representation of the problem at hand and a module that models declarative memory. All modules have parameters associated with them that determine the specifics of their behavior, such as the latency or accuracy of their operations. For example one parameter determines the

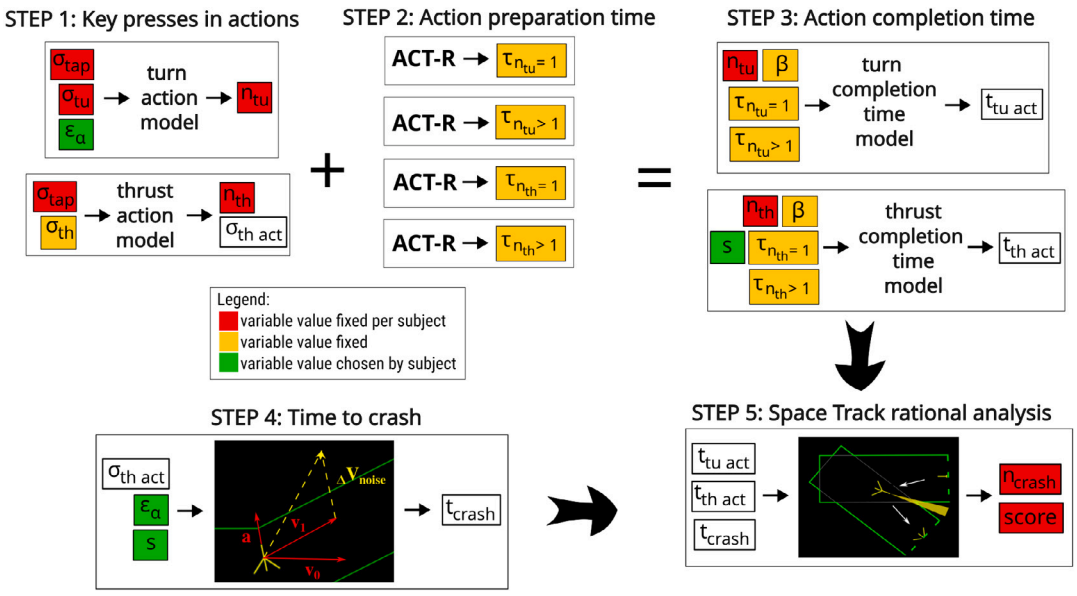

Fig. 3. Overview of tight resource-rational analysis of Space Track. We start by developing action models that predict the number of key presses needed when turning and thrusting (STEP 1). We proceed with estimating action preparation time with the cognitive architecture ACT-R (STEP 2). These models are then used to estimate turn and thrust completion time (STEP 3). Independently, we also estimate the time to crash (STEP 4). Action completion time and time to crash are the main input into the rational model of Space Track, which finds the parameter values that optimize score. For each model, we highlight the input and output variables. The variables that were fixed in independent micro-tasks are highlighted in red, those that are fixed by relying on ACT-R's default parameters are highlighted in yellow, whereas the parameters that were optimized by model and subject are in green.

rate at which long-term memory decays, whereas another determines the noise in motor actions.

Here we will rely mainly on the latency parameters of the architecture in order to estimate the time costs related to various processes. Of interest are mainly the latency parameters associated with peripheral processes (i.e., visual perception and finger movement on a keyboard) and with the procedural module. Specifically, the time it takes for a production rule to fire is fixed at 50 ms, whereas the time it takes to encode a visual object already attended to is 85 ms. The time it takes to prepare a finger movement depends on the number of features that need to be prepared. The models we construct specify the sequence in which cognitive processes are executed and produce time estimates as a sum of such latencies.

4.1.1. Step 1: Models of key presses in actions

Both turn and thrust actions are completed through either taps or long key presses. The models for both actions will rely on the tap duration and variability estimated in the Tap task. Moreover, the long turn variability estimated from the Long turn task will be input into the turn action model. When model parameters are not measured in micro-tasks, we will rely on default values from the cognitive architecture ACT-R. One such parameter common to both action models is the lower boundary of the duration of a long key press, which is 130 ms in ACT-R.[8]

[8] Motor actions in ACT-R go through a preparation phase, an initiation phase, and a burst phase. Both key press and key release will go through these 3 phases. Each phase has timing parameters associated with it. Preparation time varies as a function of the number of features that need to be prepared. The default motor initiation time is 50 ms as is the default motor burst time. Moreover, ACT-R assumes that the key press is not detected after the finger completes the movement, but already while the key travels a certain distance. The default value for this parameter, key closure time, is 10 ms. Similarly, a key release is detected before the release action is complete, specifically after a key release time of 40 ms by default. To estimate the minimum long key press, we should consider that the key release can be prepared as long as the key press is no longer in preparation. However, the key release can only start when the burst is complete. This results in minimum long key press of (burst time - key closure time) + initiation time + key release time.

4.1.2. Key presses in turn action

We consider a turn action complete when the ship has been oriented to within an aim tolerance region ε_α of a target orientation. When aiming the ship, one would typically start with a long key press (unless the orientation change is very small). In the turn action task, the first key press is expected to be long based on the turn magnitude. If that long press is unbiased (i.e., the subject aims exactly at the target orientation), the probability of falling within the aim tolerance region equals the proportion of the long turn distribution that falls within it (see Fig. C.5 in Appendix B, for a visualization of the overlap between the target orientation regions in the Turn action task and three turn distributions spanning the full range of σ_{tu}).

$$\hat{P}(n_{tu}=1;\sigma_{tu},\varepsilon_\alpha)=\int_{-\varepsilon_\alpha}^{\varepsilon_\alpha}\mathcal{N}(0,\sigma_{tu}^2)\mathrm{d}\alpha, \quad (1)$$

where $\mathcal{N}(0,\sigma_{tu}^2)$ is the density function of turn angle. The predictions of this equation for accuracies of 6 deg and 12 deg match well the empirically observed relative frequencies (Fig. 4, a) with correlations of $r_{\varepsilon_\alpha=6}=0.75$ and $r_{\varepsilon_\alpha=12}=0.89$ and little bias. Note that this equation allows us to predict the probability for any aim accuracy.

If the first key press does not fall within the aim tolerance region, corrective key presses need to be executed. These key presses can be either long presses or taps. We tested three models to investigate their nature. Model 1 assumed that all corrective key presses were long (i.e., with a SD of σ_{tu}). Model 2 assumed that they are taps, and hence were less variable than the first key press (as measured in the Tap task, σ_{tap}). This is a reasonable assumption given that the ship should be close to the target aim region after the first turn. Model 3 extended Model 2 by executing whichever (tap or long key press) gets the ship closer to the target (Appendix D presents a mathematical formulation of each model, derivations of their predictions, and a comparison of their predictions to data).

In line with Models 2 and 3, the observed duration of corrective key presses was close to the empirically measured tap averages (Fig. 4, b), albeit a little bit longer. Among these, Model 3 did best at predicting the average number of key presses in the Turn action task (with correlations $r_{\varepsilon_\alpha=6}=0.69$ and $r_{\varepsilon_\alpha=12}=0.78$, Fig. 4, c). To summarize, the turn model assumes that subjects orient a ship to a target accuracy

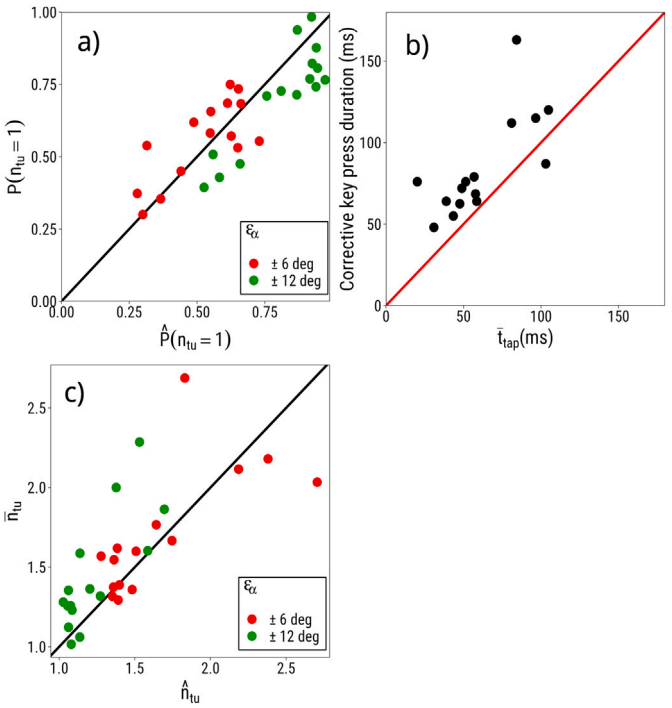

Fig. 4. Turn action model plots. (a) Predicted and observed (Turn action task) probability of completing the turn in a single key press. (b) Tap duration (Tap task) and median corrective key press duration (Turn action task). (c) Predicted and observed (Turn action task) number of key presses in a turn action.

by first executing a long key press, possibly followed by corrective key presses, which are in most cases short taps (accuracy as measured in the Tap task), but also long key presses on occasion.

4.1.3. Key presses in thrust action

The goal of a thrust action is to change the ship initial velocity to its target velocity. Before the thrust action is initiated, the ship is aimed appropriately such that the subsequent thrust can reach the target velocity. We consider a thrust action complete if the ship cannot get any closer to that velocity without reaiming, that is, if the velocity change produced by the thrust is at most half a tap below the target velocity change. If the velocity change is further below, the model will execute appropriate corrective thrust key presses.

Unlike for turn key presses, we have not measured long thrust variability σ_{th}. Instead, we will use a common parameter for all subjects. Just as for long turn key presses, we will assume that the variability of long thrust key presses is independent of press duration. Moreover, we will assume that the same coefficient of variation found for taps (i.e., 0.35, see Fig. C.1 in Appendix B) also applies for long thrusts and estimate its SD as that for the smallest long key press: $\sigma_{th} = 0.35 \cdot 130 = 45.5$ ms. The thrust action model assumes that an unbiased long key press is executed with that SD, followed by one or several corrections if the first thrust was shorter than the target thrust. The number of corrective key presses is a function of how far the first thrust is below the target thrust. For example, if it is between 0.5 and 1.5 taps below the target thrust, a single tap will bring the ship closest to its target. If it is between 1.5 and 2.5 taps below, 2 taps will be necessary, and so on. The final estimate of expected key press number in a thrust action $\hat{n}_{th}$ weighs the number of corrective taps by the probability that this number of taps will be needed. This estimate matches well the observed number of key presses in a thrust action in Space Track $\bar{n}_{th}$ ($r = 0.78$, see Fig. 5). Note that thrust corrections can be executed after the initial thrust, which means that the thrust action variability $\sigma_{th\ act}$ will be smaller than that of a long thrust σ_{th}.

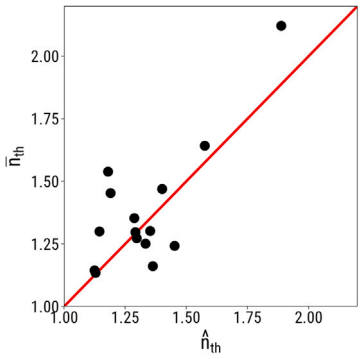

Fig. 5. Predicted and empirical average number of key presses in a thrust action. The empirical average number of key presses is estimated in thrust actions in the main task, Space Track.

Table 4
Estimated time costs associated with the cognitive processes responsible for triggering an action and their corresponding game values. Empirical Turn correction cognitive costs were computed as the time between consecutive turn key presses. The median inter-key-press interval of the median subject is presented.

Action	Notation	ACT-R time estimate	Data (median)
First turn	$\tau_{n_{tu}=1}$	570 ms	N/A
Turn correction	$\tau_{n_{tu}>1}$	335 ms	300 ms
First thrust	$\tau_{n_{th}=1}$	285 ms	N/A
Thrust correction	$\tau_{n_{th}>1}$	285 ms	267 ms

4.2. Step 2: Cognitive delay

Action completion time t_{act} is the sum of turn completion time $t_{tu\ act}$ and thrust completion time $t_{th\ act}$. Each of these consists of a cognitive delay τ (i.e., the time costs of all cognitive processes that prepare the action) and key press time associated with potentially several key presses. Here we estimate the cognitive delay τ by constructing cognitive models of these actions in ACT-R.

We provide detailed process traces of the cognitive models in Appendix E and their time cost estimates in Table 4. To summarize these models, preparing the first turn (i.e., $n_{tu} = 1$) requires more work than turn corrections (i.e., $n_{tu} > 1$). The subject needs to first attend to the inter-segment angle β and extract its bisector in order to orient the ship along it. Once extracted, the subject needs to attend to the ship and start either a clockwise or a counterclockwise turn. These processes are estimated to take $\tau_{n_{tu=1}} = 570$ ms. Preparing turn corrections is much simpler, because the subject has already determined the bisector orientation and is already attending to the ship. The subject needs to merely perceive its orientation and initiate the next key press, which is estimated to take $\tau_{n_{tu>1}} = 335$ ms. Finally, preparing thrusts requires the same processes as turn corrections, but the motor command can be prepared in advance, since it is always the same. This saves an additional 50 ms for $\tau_{n_{th=1}} = \tau_{n_{th>1}} = 285$ ms. The time estimates of preparing a correction are close to subjects' median (see Table 4; estimated through the median inter-key-press-interval for all consecutive turn key presses and thrust key presses in the main task).[9]

[9] Some might argue that we can further validate our model predictions from the Turn action task. However, inter-key-press-intervals in the Turn action task are much longer than in the Space Track, because subjects have to attend to visual feedback on whether their turn was successful or not. Hence, they are not good estimates of turn correction cognitive delay. Moreover, in Space Track there might be a waiting period between completing a turn and starting the next thrust, and between completing a thrust and starting the next turn, which make it hard to generate a reasonable empirical estimate of cognitive delays related to the first turn or thrust.

4.3. Step 3: Action completion time in Space Track

Turn and thrust completion times consists of cognitive delay and key press time. Key press time can be estimated based on target value and game parameters (i.e., the two constants: ship acceleration a, which specifies by how much velocity changes each game tick when the thrust key is pressed, and turn rate ω, which specifies by how much orientation changes each game tick when a turn key is pressed). Specifically, thrust key press time equals the absolute change in velocity divided by ship acceleration $|\Delta v|/a$, whereas turn time equals the change in angle divided by turn rate $\Delta\alpha/\omega$. Note that velocity change is a vector and here we take its absolute value to compute the vector length. Also note that key press time will vary around the ideal value as a function of subject tap and long key press variability.

Our model assumes the same cognitive delays τ for all subjects. However, it does predict different action completion times based on the number of key presses subjects need to complete turn and thrust actions. Turn completion time is most strongly affected by the number of key presses n_{tu} a subject needs to reach an aim accuracy ε_α. All subjects will need the same amount of time to prepare and execute the first key press, but inaccurate subjects (i.e., those with a high σ_{tu} and σ_{tap}) will need to prepare and execute one or a few corrective key presses as well, especially at low aim tolerances (Fig. 6, a; green is the turn completion time distribution of an inaccurate subject, whereas red is of an accurate one). For example, when $\varepsilon_\alpha = \pm 6$ deg, an accurate subject will complete most of her turns within 2 s, whereas an inaccurate subject will require as long as 4 s on occasion.[10]

Thrust completion time depends on tap duration t_{tap}, speed s and inter-segment angle β (Fig. 6, b). A faster flying ship will thrust longer to steer in a new direction. Moreover, at a large β the ship needs to almost reverse its flight direction, a larger velocity change, which requires more time than the more gentle adjustments required at small β. Finally, subjects with very short taps will need to execute more key presses n_{th} to correct too short thrusts. Overall, thrust completion time can be as short as 0.5 s when target speed is slow and no corrective key presses are needed and as long as 2.5 s when target speed is fast and corrections are executed. Typically, turn completion time is longer than thrust completion time, especially for an inaccurate subject flying slowly. However, for a fast flying accurate subject, these two action types take a similar amount of time to complete.

4.4. Step 4: Time to crash

Time to crash t_{crash} is a function of the speed s at which the ship flies and the distance to segment border d_θ in its flight direction θ. Since a segment is longer than wide, if the ship is heading toward the segment side, d_θ will typically be short and, consequently, t_{crash} will be small. After an inter-segment transition, the ship would ideally fly along the second segment at target speed s_{tar} (red velocity vector v_1, Fig. 7). The ideal velocity would be disturbed by imperfections in subject actions. Specifically, the amount of noise in turn (i.e., ε_α) and thrust actions (i.e., $\sigma_{\text{th act}}$) would determine velocity noise Δv_{noise}. If Δv_{noise} is large enough, t_{crash} might be strongly affected, either because the ship is heading to the segment side or because its speed is much faster than intended.

How exactly velocity noise influences time to crash t_{crash} will vary with ship speed s and inter-segment angle β. Specifically, ship speed s modulates the relative influence of Δv_{noise} on flight direction θ: the faster the ship flies, the less will velocity noise Δv_{noise} distort its flight direction. Inter-segment angle β affects the angle between ship aim and target flight direction. At large β, ship aim is close to perpendicular to the target flight direction (Fig. 7, a and b), whereas at small β, ship aim is close to parallel (Fig. 7, c and d). In the first case (Fig. 7, a), Δv_{noise} will disturb ship flight direction θ a lot and resultant speed to a lesser extend, whereas in the second case (Fig. 7, c) it will mostly affect the resultant speed and disturb flight direction to a much lesser extend.

If the ship is not ideally aimed (Fig. 7, b and d), then its speed will have to be different than the ideal one if the subject aims to achieve a flight direction along the second segment. How large the speed disturbance Δs will be will vary with β. When β is large, aim noise will minimally affect resultant speed (Fig. 7, b). When β is small, compensating for aim noise might require large adjustments in target speed to achieve a flight direction along the second segment (Fig. 7, d).

These effects will mean than at small β, mostly aim accuracy ε_α will influence if the ship is heading to the segment back or side (Fig. 8) and, consequently, what t_{crash} will be (Fig. 9). Here, when ship aim is inaccurate, the ship might crash into the segment side in less than 1 s after thrust is complete, even at moderate speeds. At large β, the effect of thrust noise $\sigma_{\text{th act}}$ will predominate for both variables and t_{crash} will be a little longer (>2 s), even when headed to the side.

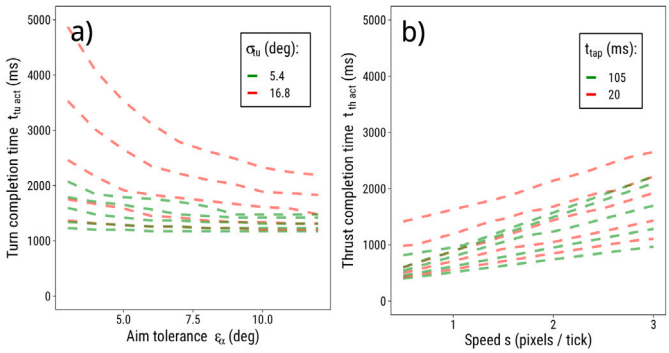

Fig. 6. Predicted turn completion time (panel a) and thrust completion time (panel b). The lines represent the 10th, 25th, 50th, 75th and 90th percentiles. We plot the predicted distribution of a low completion time subject (green) and a high completion time subject (red) for both turn and thrust actions. The distributions differ, because the high completion time subject needs more key presses, on average, than the low completion time one. The time distributions are a function of both the distributions over the number of key presses to complete the action and various in-game parameters, such as how far the ship turns and how long it thrust for.

4.5. Step 5: Putting it together

The complete model estimates the probability that the ship crashes as a function of the parameters discussed in the previous sections. There are many opportunities for the ship to crash in a segment if the subject does not act fast enough, either during a trajectory correction (of which there can be several in a segment) or during the inter-segment transition. We label the cumulative probability of those *the segment crash probability* p_{c} (please refer to Table 5 for a list of all variables used in this section). Note that p_{c} is a function of many variables, such as how the turn action in the main inter-segment transition was executed, how thrust actions in trajectory corrections are executed, or what trajectory the ship had followed in the previous transition.

If the ship crashes, all its progress is lost and it reappears at the beginning of the segment, only to potentially crash again. The ship will continue to reappear at the beginning of the segment after each crash until it enters the subsequent segment. This means that the number of crashes before successfully traversing a segment n_{c} can be described by a geometric distribution, whose average $\bar{n}_{\text{c}}$ is a function of the probability of crashing on each traversal attempt:

$$\bar{n}_{\text{c}} = \frac{p_{\text{c}}}{1 - p_{\text{c}}}. \tag{2}$$

[10] Note that we assume that trajectory corrections rely on a highly inaccurate aim and, hence, there aiming the ship is always completed in a single key press. This is because trajectory corrections are equivalent to inter-segment transitions at small β, which are insensitive to aim accuracy.

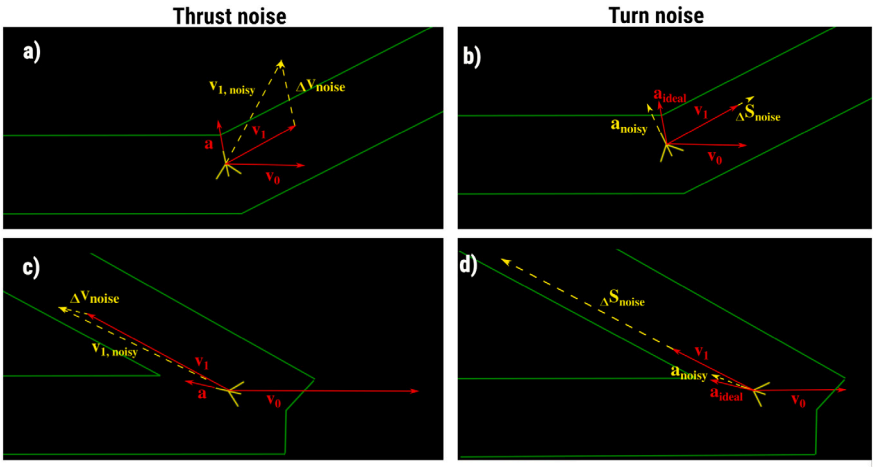

Fig. 7. Affect of thrust and turn noise on resultant speed for small and large inter-segment angles. Red vectors depict the ideal scenario, whereas yellow vectors the effect of noise. (a) Affect of thrust noise at a large inter-segment angle. Here thrusting is close to perpendicular to flight direction and therefore thrust noise leads to a relative large change Δv_{noise} in resultant velocity v_1. This change affects flight direction a lot when target speed is small. (b) Affect of turn noise at a large inter-segment angle. When the ship thrusts a few degrees form its ideal thrust direction, the ship's final velocity is just barely affected. (c) Affect of thrust noise at a small inter-segment angle. Thrust direction is close to parallel to the target flight direction, leading to minor influence of thrust noise on resultant flight direction. This influence is even smaller when speed is large. (d) Affect of turn noise at a small inter-segment angle. Here minor deviations from the target thrust angle can lead to large changes of resultant speed to reach the target flight directions.

Table 5
Estimated time costs associated with the cognitive processes responsible for triggering an action and their corresponding game values. Empirical Turn correction cognitive costs were computed as the time between consecutive turn key presses. The median inter-key-press interval of the median subject is presented.

Variable	Notation
Segment crash probability	p_c
Segment crash frequency	n_c
Successful segment traversal time	T_s
Segment crash time	T_c
Segment time	T_{seg}
Number of traversed segments in a game	n_{seg}
Crash frequency in a game	n_{crash}

Each crash naturally increases the time the ship spends in a segment because progress in the segment is lost and, also, the ship only reappears 1 s after the crash. Given the ship's speed, it will traverse a segment successfully in T_s seconds. If the ship fails to traverse a segment because of a crash, it will take on average T_c seconds. The reset cost of 1 s needs to be added to this to compute the overall time cost of a crash. The average time the ship needs to traverse a segment T_{seg} will include both the average time lost because of crashing and the time to successfully traverse the segment:

$$T_{seg} = T_s + \bar{n}_c(T_c + 1). \tag{3}$$

If a subject needs T_{seg} seconds to traverse a segment, on average, s/he will traverse $n_{seg} = 180/T_{seg}$ segments in a single 3-minute game. Game score will weight traversals by 25 points and crashes by −100 points, $25 \cdot n_{seg} - 100 \cdot n_{crash}$, where $n_{crash} = n_{seg} \cdot n_c$.

These parameters are all related to the parameters estimated in the previous section. Specifically, p_c is the probability that $t_{crash} < t_{act}$ for the inter-segment transition and all trajectory corrections. T_s is the time the ship needs to traverse the segment given its velocity and β. This time is a little bit longer than it would be if the ship were flying at its target velocity at all times, because the ship temporarily slows down during inter-segment transitions and this slowing down is more significant the sharper the angle. Finally, T_c is the expected value of t_{crash} for all t_{crash} shorter than t_{act}. To compute these values, we will exploit the symmetry of the task.

To exploit the task's symmetry, we will introduce the notion of periodic boundary conditions. Periodic boundary conditions are used in physics to describe symmetric large systems, such as crystals or liquids. Specifically, when assuming periodic boundary conditions, one models the entire system with a small part of it (the so called *unit cell*) that repeats infinitely to produce the large system. For example, the electric potential in a crystal would be approximated by the potential around a single atom or collection of few atoms that constitute the smallest unique part of the structure (i.e., the unit cell). An electron would be assumed to reappear on the opposite side of the unit cell once it exits from one side. Analogically, we can model a track in space, in which all transitions have the same inter-segment angle, as a single inter-segment transition consisting of two intersecting half-segments (see Fig. 10). In this setup, the ship reappears at the beginning of the first half-segment after it exits the second one. This reduces a large track of multiple transitions to a single transition.

Assume that the ship always enters the unit cell at the same speed s_{tar}, that it only sees a single β, that the ship is always aimed to within the same aim accuracy ε_α, and that the thrust action is always executed at the ideal moment with the same velocity noise. The inter-segment transition would take t_{act} seconds to complete and would result in a time to crash distribution t_{crash}, and the ship would crash whenever $t_{act} > t_{crash}$. Then an optimal model would choose s_{tar} and ε_α which produce the largest score. This setup demonstrates that smaller β are more difficult than larger ones, because they require longer thrusts to almost reverse the ship flight direction and more turn key presses to reach narrow aim tolerance regions that they require. Average score at the lowest β is consistently 1000 points less than at large β.

For the complete model, we relax two assumptions of this initial setup. First, we introduce *probabilistic periodic boundary conditions* by assuming that each β is equally likely instead of assuming a single β. The model needs to choose a single value for s_{tar} and ε_α that works well across all β. Because crashes are more costly than segment traversals, this model will be biased towards values that work well at small β, that is narrow ε_α and slow s_{tar}. Second, we will relax the assumption that the ship always enters at its ideal speed s_{tar}. Instead, we will assume that the speed at which the ship enters the unit cell is distributed just as the speed at which it exits the unit cell if it does not crash. Often the ship will fly at close to the ideal speed, but also far from it on occasion. The further the ship is from the optimal speed, the lower the expected score will be. Since less capable subjects will fly the ship at a more variable speed, their performance will be more strongly penalized by this assumption.

Optimal s_{tar} and ε_α values were found using grid search. The resultant model predictions correlated strongly with subject asymptotic

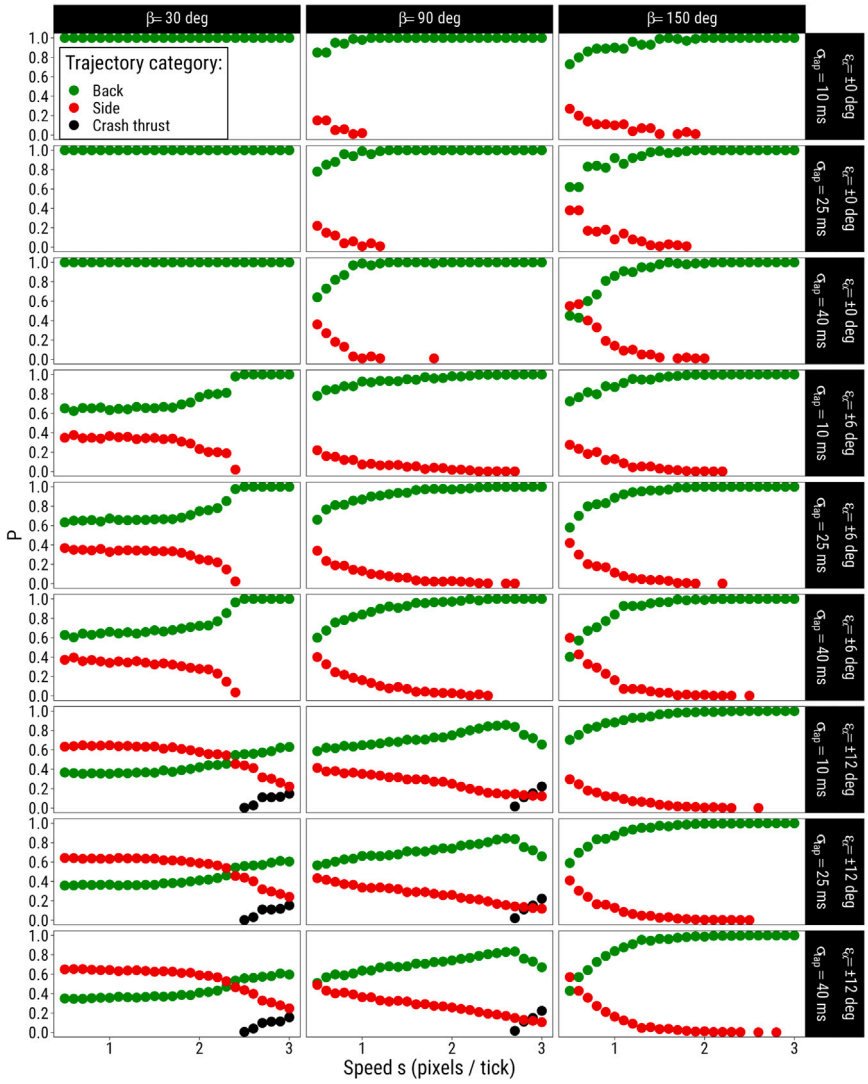

Fig. 8. Probability of heading toward the back or side of the segment, or crashing before completing thrust, as a function of speed s, tap noise σ_{tap}, and aim tolerance ϵ_α. At $\epsilon_\alpha = 0$, the ship is aimed at exactly the target orientation. At $\beta = 150$ deg, aim tolerance barely influences flight direction, whereas at $\beta = 30$ deg it has a strong influence. Conversely, thrust noise strongly affects where the ship will head at large β, whereas at small β its affect is much more modest. Note that when $\epsilon_\alpha = 0$, the ship is always aimed at the inter-segment angle bisector.

game score ($r = 0.87$) and game crashes ($r = 0.86$, Fig. 11) and were generally close to observations in absolute terms (mean absolute error: MAE(score) = 332; MAE(n_{crash}) = 1.89). Moreover, predicted score was, on average, slightly higher than observed score (mean signed deviation: MSD(score) = 163), whereas predicted crashes were close to observations (MSD(n_{crash}) = 0.15). The conclusions were drawn with asymptotic performance estimated over a window of 4 games, yet they remain robust when we systematically vary the window size (for a study of the effect of window size over these measures, see Appendix G). The model predictions are generated entirely from parameters measured in micro-tasks, without estimating any parameter values from the main task. The majority of subjects are very close to their predicted performance. Moreover, the model quantitatively predicts that most subjects will rarely crash, whereas a few will crash much more frequently.

5. General discussion

We developed a resource-rational model of Space Track, a task in which producing rapid and precise actions is key to success. Importantly, both the time cost and the accuracy of actions are a function of multiple cognitive processes. We developed models of these actions, which specified action accuracy and latency. To estimate action accuracy, we created mathematical models, whose parameters were fixed per subject in independent micro-tasks. The models were tested either in an independent micro-task or in the main task. To estimate action latency, we constructed ACT-R models and relied on ACT-R's default parameters. After propagating the effects of action accuracy and timing onto the main task and assuming that each subject did their best, we arrived at individualized score predictions close to asymptotic score. This provided strong support for our model. Importantly we demonstrated a method of estimating resource-rational model parameters in independent micro-tasks that is useful in constraining complex models. We will now discuss how and why our model deviates from the data and then focus on different aspects of our modeling approach.

5.1. Deviations of model predictions from data

The predictions of our model corresponded well to observations. Thus, we demonstrated that perceptual-motor limitations have a strong influence on Space Track performance. Yet, even though model predictions and data were close, model predictions were above observations

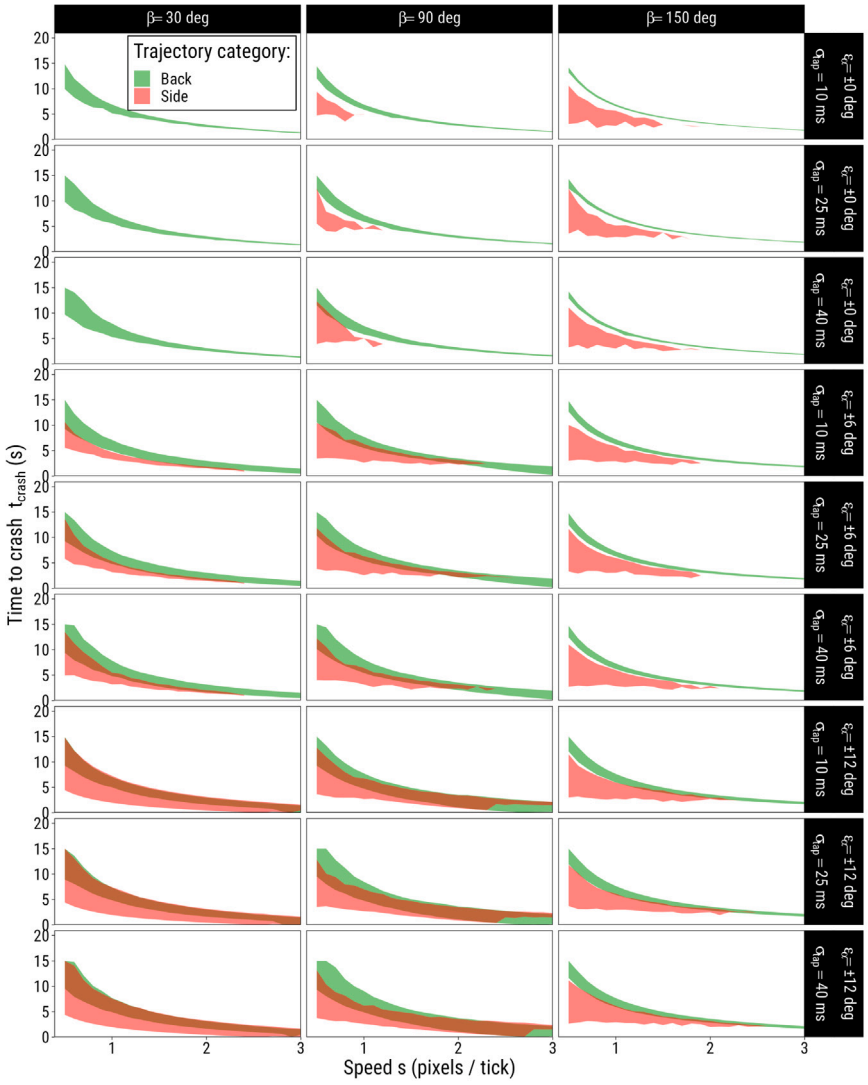

Fig. 9. Time to crash t_{crash} after an inter-segment transition as a function of speed s, tap noise σ_{tap}, and aim tolerance ε_α and flight direction category. At $\varepsilon_\alpha = 0$, the ship is aimed at exactly the target orientation. Plotted are regions between 5th and 95th percentiles of t_{crash}.

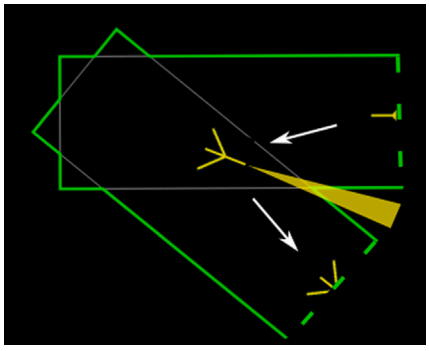

Fig. 10. A depiction of the assumption of periodic boundary conditions in Space Track. In this setup, the ship faces the same inter-segment angle. When the ship exits the unit cell, which consists of two half-segments at the corresponding inter-segment angle, it reappears at the start of the first half-segment. The ship is assumed to always select the same aim accuracy (highlighted in yellow), which in turn determines the turn completion time distribution. Moreover, the ship picks a single target speed to fly at, which determines the thrust completion time distribution and the time before crash. These jointly determine the probability of crashing and the number of segments that will be traversed, and consequently, final score.

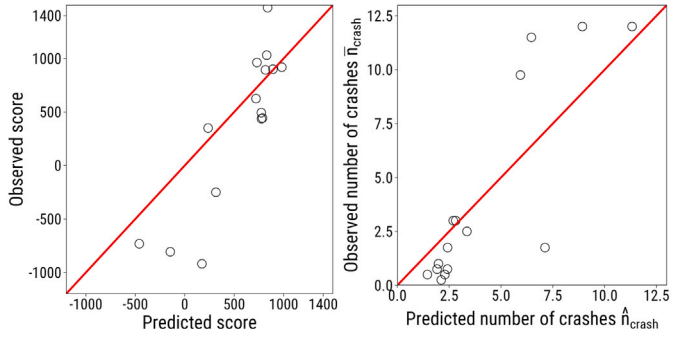

Fig. 11. Predicted asymptotic Space Track score and crash frequency in segment vs. empirically observed values in the best performing 4-consecutive games of study. The line of equivalence is plotted in red.

for many subjects. Moreover, the model significantly overpredicted one poor performing subject and underpredicted the best performing subject. The discrepancy between predictions and observations can come from violations of the assumption of rationality, from how Space Track score is presented to subjects, from other measurement errors, or from assumptions in the model. We will now look at each of these separately.

Humans do not always behave rationally, because our learning processes do not always find the optimal solution (e.g., Fu & Gray, 2006). Even in the original experiments with Space Fortress (e.g., Donchin, 1995), a game on which Space Track is based, subjects were content with finding any strategy that produced positive score, whether optimal or not. Generally, in complex tasks, the space of possible solutions is very large and subjects might not always find the global maximum, at least within the time course of an experiment. In fact, skill acquisition models of Space Track (Anderson et al., 2019; Seow et al., 2021) have demonstrated that learning processes can introduce suboptimalities in performance.

However, it is difficult to claim that subjects with a negative score potential are behaving suboptimally when scoring below their estimated potential. While we relied on the true Space Track score in our analysis, this was not the score that subjects saw in the experiment. When designing the study, we set a lower limit on score of 0 to avoid frustrating poor subjects.[11] Consequently, from subjects' point of view, all negative score was equally good. This might also explain why one subject scored worse than a passive subject and others started decreasing their score after a few games, which further tainted our measurements.

In fact, this is not the only possible measurement error in our approach. Previous experiments with Space Track have shown an increase in score even after 40 games, whereas subjects only played 20 games in our study. Thus, it might be that subject were not given enough time to reach their potential. We made this choice for practical reasons as previous experiments have also shown that most learning is complete within 20 games. In fact, providing subjects with raw scores might have given valuable additional feedback that would speed up learning.

In addition to these data-centric problems, certain assumptions in our model might be the cause for its predictions error. A major simplification in the model is that it only relies on individual measures related to action accuracy. For all other parameters, we used values of an average subject from ACT-R. Individualizing all action parameters might have produced a more accurate prediction. In fact, this is our best guess about why the best subject is scoring higher than what our model predicts: that subject likely payed smaller cognitive preparation costs than ACT-R's default parameters. Another simplification in the model was to treat the setup as that of periodic boundary conditions with a distribution of starting speeds around the optimal. This idealization might have been too coarse for the poor performing subjects. Despite these potential pitfalls of our model, we do find it useful for two reasons. First, it demonstrates that subjects can be close to optimal even in tasks such as Space Track and, consequently, that learning processes would likely introduce only minor suboptimalities in tasks of similar complexity. Second, it showcases that tasks of such complexity are also amenable to a rational analysis.

5.2. Comparison with other approaches

Our approach naturally comes closest to other resource-rational analyses, because of its focus on the impact of cognitive constraints on task performance. It is perhaps closest to that of Howes et al. (2009). Just as we do, they incorporated a cognitive architecture into a rational analysis and, also, they calibrated individual model parameters. The difference is that they did not calibrate the parameters in independent micro-tasks, but on the same task. Another difference is that they did not need to validate model components in independent tasks. The latter might be specific to more complex tasks, which require multiple different component skills, each of which requires developing an independent model.

In this sense, our approach is also close to certain skill acquisition approaches. Skill acquisition models and rational models are generally complementary, because the former aim to characterize learning processes, whereas the latter aim to characterize behavior once learning has done its job. Yet, Hoekstra, Martens, and Taatgen (2020)'s skill-based approach has interesting touching points with our approach. They emphasize the flexibility that humans exhibit in reusing skills across multiple tasks and construct models that exhibit a similar amount of flexibility. At the core of their modeling approach are multiple simple skills that are shared across models. This is analogical to the turning and thrusting models that we constructed and validated as components of the Space Track model. The commonality between the two approaches, we believe, is componential decomposition, which is a natural step to take when dealing with complexity.

5.3. Tight resource-rational analysis

We offer tight resource-rational analysis as a method of dealing with model complexity. This method is complementary to other methodological approaches, which focus on penalizing complex models either through an explicit penalty on the number of free parameters (e.g., Akaike Information Criterion or Bayesian Information Criterion) or an implicit penalty based on the size of the space of all possible predictions (e.g., cross-validation). Our approach reduces model complexity by fixing parameter values either through individualized measurement in separate micro-tasks or through using group-level average values from an established unified theory of cognition. Moreover, our approach reduces possible variation in functional form by validating functional form hypotheses in additional micro-tasks. The thus constrained models only vary within the bounds of parameter measurement error, which, we believe, increases confidence in their validity.

It is apparent why fixing parameters in independent micro-tasks is not often performed: it is time consuming.[12] Even with only three micro-tasks, the duration of our experiment roughly doubled. Some tasks might require tens of accompanying micro-tasks to constrain model parameters, which would be difficult to implement. In fact, even in our experiment we might have needed many more micro-tasks had we not relied on ACT-R. The key is being selective and focusing on parameters that would account for most of the inter-subject variability in performance. In our case, we made an informed guess about which parameters are most likely to be central to the task by the relying on the existing ACT-R model and our intuition about what would affect the probability of crashing in a segment most. In essence, one needs a rough sketch of the detailed resource-rational model when deciding what to measure. Despite this drawback, we believe that any researcher that is moderately familiar with a task can make a informed guess about what to measure and what not. Overall, we consider our approach an important methodological stepping stone towards implementing resource-rational models of more complex tasks.

CRediT authorship contribution statement

Cvetomir M. Dimov: Writing – review & editing, Writing – original draft, Visualization, Methodology, Formal analysis, Data curation, Conceptualization. **John R. Anderson:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Funding acquisition. **Shawn A. Betts:** Software.

[11] On occasion, a true negative score can also be seen as a positive score if the subject was performing relatively poorly at the beginning of the game, but relatively well toward the end.

[12] As the great tradition of measuring individual differences in psychometry has demonstrated (e.g., Ackerman, 1987).

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: John Anderson reports financial support was provided by Air Force Office of Scientific Research.

Data availability

All data and data analysis code can be found on the Open Science Foundation website at 10.17605/OSF.IO/MGFXH.

Appendix A. Micro-task rewards

Performance in all micro-tasks was incentivized through a quadratic reward function. The specific values of the reward was calibrated such that the average subject would achieve a reasonable number of points.

A.1. Tap task

The goal of this task was to measure the shortest key press subjects could execute. Consequently, shorter key presses were awarded more points. The specific reward function R was chosen such that only key presses of below 100 ms are awarded any points. Moreover, the denominator was chosen such that a key press of a duration that half of all subjects can achieve (i.e., 50 ms) would bring 1 point.

$$R = \begin{cases} \frac{(100-t)^2}{2500} & \text{t} < 100 \text{ ms} \\ 0 & \text{else} \end{cases}$$

A.2. Long turn task

The goal of this task was to measure the variability of a long turn key press. Only turns that fell within 15 degrees of the target orientation were awarded any points. Moreover, turns that were 8 degrees away would receive 1 point, with the theoretically highest award being 4 points, when the offset is 0 degrees.

$$R = \begin{cases} a\frac{(15-\Delta\alpha)^2}{64} & \Delta\alpha < 15 \text{ deg} \\ 0 & \text{else} \end{cases}$$

A.3. Turn action task

The goal of this task was to measure how quickly a subject can orient the ship to a target region. Consequently, rapid action was incentivized. All trials that took longer than 1500 ms did not bring any points. Moreover, a trial that would take 707 ms would bring 1 point. Note that timing was initiated when the first key press was detected and not at the very start of the trial.

$$R = \begin{cases} \frac{(1500-t_{tu\ act})^2}{500000} & t_{tu\ act} < 1500 \text{ ms} \\ 0 & \text{else} \end{cases}$$

Appendix B. Data cleaning

We are interested in measuring asymptotic performance in all tasks. To this end, we eliminated all learning trials on all tasks. For micro-tasks, we removed the first 25 trials and for Space Track, we only considered the 4 games with highest average performance. Additional data cleaning was performed for the Tap and Long turn tasks. First, some subjects continued improving their performance on a task or did not reach a stable-state performance until later. For those subjects later trials were considered. Second, outliers were eliminated by not considering all data points that lay outside of 1.96 standard deviations around the mean. In the following subsections, we highlight which data were removed (green: learning trials; red: outliers) and which kept (black) for the all tasks.

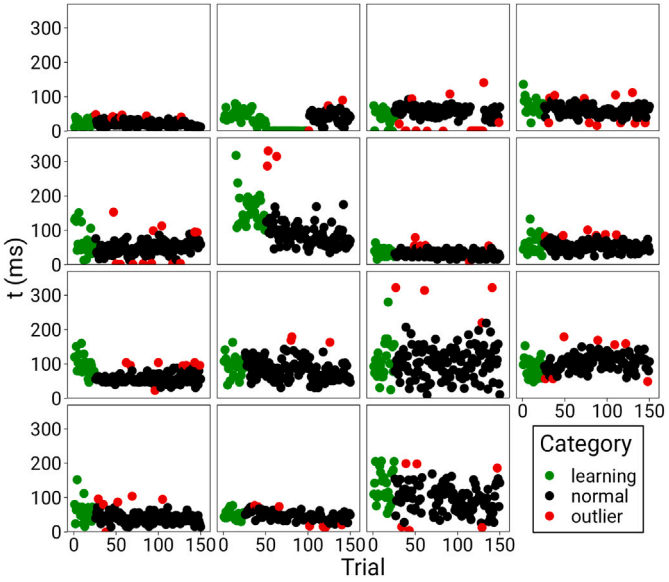

Fig. B.1. Per-subject data from Tap task. Each panel plots a subject. The first 25 trials were removed for all subjects as subjects were assumed to still be learning the task (green). Two subjects required removing more data. Among the remaining subjects, outliers were eliminated (red). Distributions were fitted to the remaining data (black).

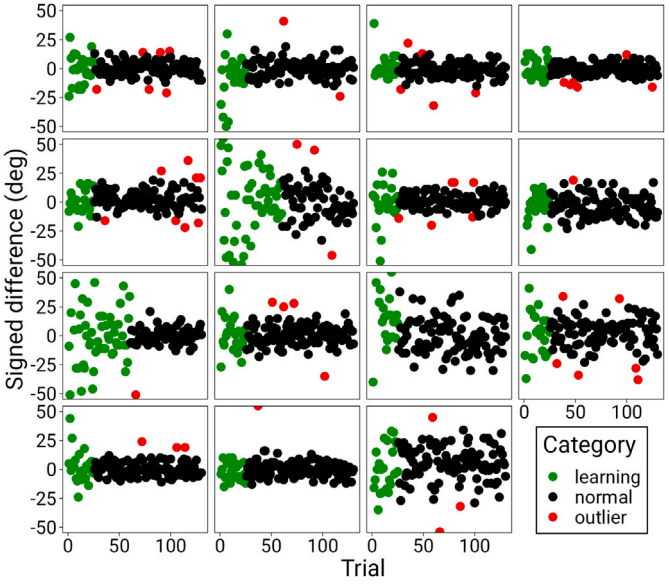

Fig. B.2. Per-subject data from Long turn task. Each panel plots a subject. The first 25 trials were removed for all subjects as subjects were assumed to still be learning the task (green). Two subjects required removing more data. Among the remaining subjects, outliers were eliminated (red). Distributions were fitted to the remaining data (black).

B.1. Tap task

Fig. B.1 plots learning, outlier and normal data for the Tap task. For the majority of subjects, the first 25 trials were considered learning trials. In addition, one subject (first row, second column) exhibited unrealistically fast response times in the middle of the task. For this subject, we extended learning trials to include all trials until the end of this fast streak (trial 100). A second subject (second row, second column) continued decreasing response time after trial 25, which also necessitated considering additional trials as learning trials. As described above, we removed the 5% most extreme data points as outliers. Before eliminating these outliers, all trials with duration of 1 or 0 ms were removed (artifacts). These are highlighted as outliers too.

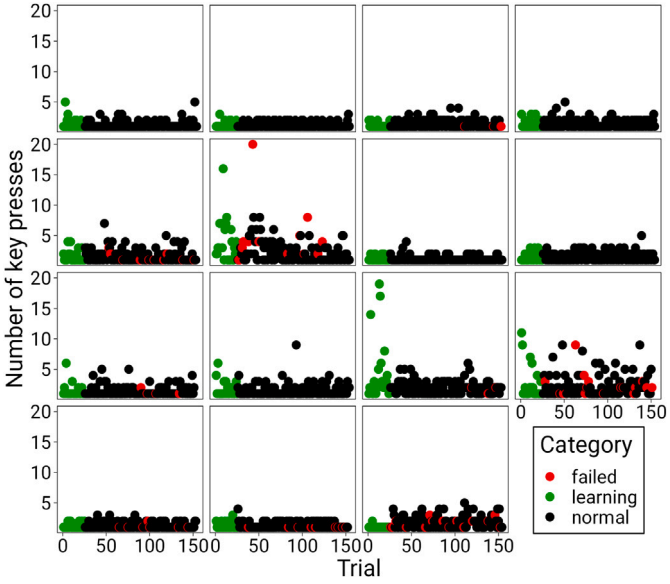

Fig. B.3. Per-subject data from Turn action task. Each panel plots a subject. The first 25 trials were removed for all subjects as subjects were assumed to still be learning the task (green). Among the remaining subjects, failed trials were eliminated (red).

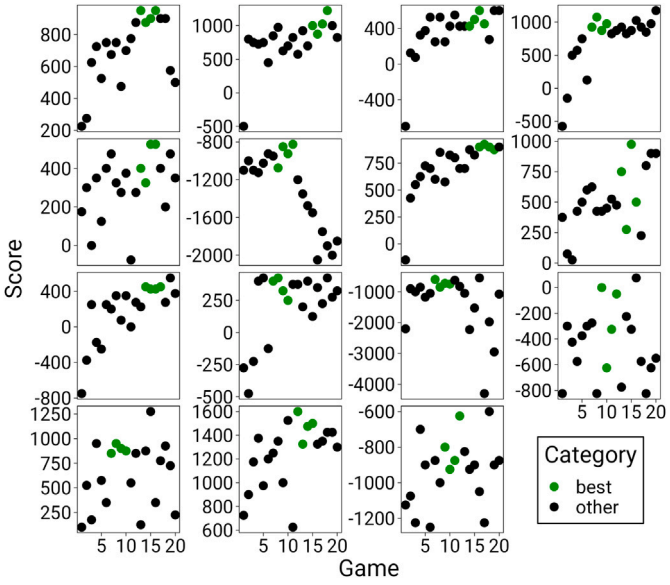

Fig. B.4. Per-subject data from Space Track. Each panel plots a subject. The four consecutive games with the highest average score were selected as an approximation of asymptotic performance.

B.2. Long turn task

Fig. B.2 highlights learning, outlier and normal trials in the Long turn task. For the majority of subjects, the first 25 trials were considered learning trials. Two subjects did not stabilize their performance until later (second row, second column; first row, third column). For these, trials up to trial 60 were considered learning trials. The 5% most extreme data points were removed as outliers.

B.3. Turn action task

Fig. B.3 highlights learning, failed and normal trials in the Turn action task. The first 25 trials were considered learning trials for all subjects. Failed trials are those, in which subjects could not orient the ship to within the target region within the trial time limit. Those were removed from further consideration.

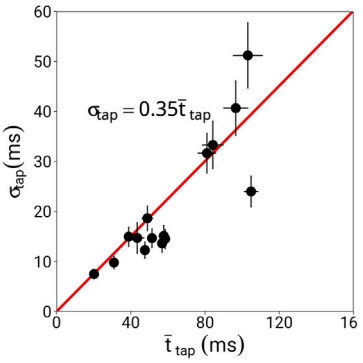

Fig. C.1. Mean tap duration and variability. Each observation corresponds to a single subject. Error bars represent confidence intervals. Tap duration exhibits a scalar variability (i.e., mean tap duration is proportional to tap SD) with a coefficient of 0.35.

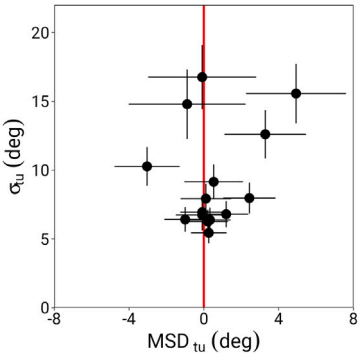

Fig. C.2. Long turn mean signed difference (MSD) and variability (SD). Each observation corresponds to a single subject. Error bars represent confidence intervals. The majority of subjects are not biased (MSD not significantly different from 0).

B.4. Space Track

To approximate subject optimal performance, we took the four consecutive games that had the highest average score (highlighted in green in Fig. B.4). We did not take the last 4 games, because some subjects showed a decrease in performance in the later games (e.g., row 1, column 1; row 2, column 2), probably due to lack of motivation.

Appendix C. Additional data analysis plots

The main text provides only summary measures of quantities of interest. Here, we present plots related to each micro-task for the interested.

C.1. Tap task

Tap duration exhibits scalar variability. The majority of subjects could execute taps shorter than 100 ms, but two subjects were much slower, both at around 200 ms. These two subjects were removed from further analyses.

C.2. Long turn task

The majority of subjects produced unbiased long turns. They varied widely in variability (SD between 5 and 19 degrees; see Fig. C.2). Unlike the previous task, no outlier subjects were detected. We also visualize the long turn variability of three example subjects: a good one, an average one, and a poor one (see Fig. C.3).

Fig. C.3. Three example turn distributions, spanning the full range of subject single turn variabilities. The red line represent the ship's target orientation. The ship's turn is assumed to be unbised. More saturated yellow regions correspond to regions of higher probability density.

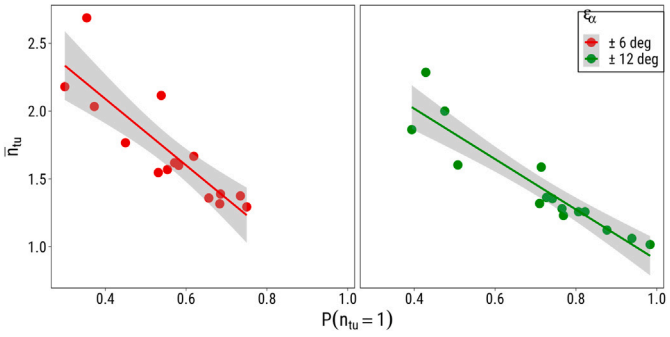

Fig. C.4. Number of key presses in a turn action and probability of completing a turn action with a single key press in the Turn action task. The two panels (and colors) represent the two tolerance regions. There is a strong linear relationship between these.

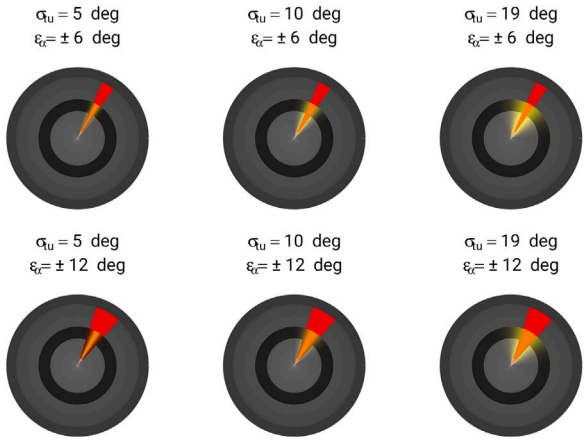

Fig. C.5. Three example turn distributions, spanning the full range of subject single turn variabilities, superimposed over the two tolerance regions in the Turn action task. The ship's aim is assumed to be unbiased (i.e., on average, it is in the center of the tolerance region). More saturated yellow regions correspond to regions of higher probability density. The proportion of the yellow curve that falls within the red region correspond to the probability of falling within the tolerance region in a single turn.

C.3. Turn action task

Subjects varied too in the average number of key presses they needed to complete a trial in the turn action task. Here we provide plots relating the two measures of interest: average number of key presses in a turn action $\bar{n}_{tu}$ and probability of completing a turn action in a single key press $P(n_{tu} = 1)$ (Fig. C.4). Similar to the long turn task, we also visualize subject long turn SD relative to the two tolerance regions in our experiment (see Fig. C.5).

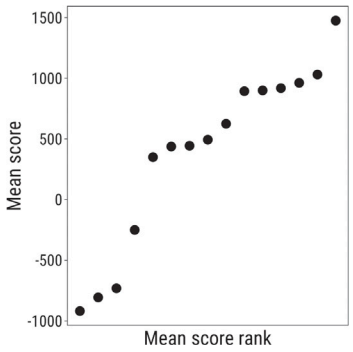

Fig. C.6. Average subject score over best performing 4 games, ranked by score.

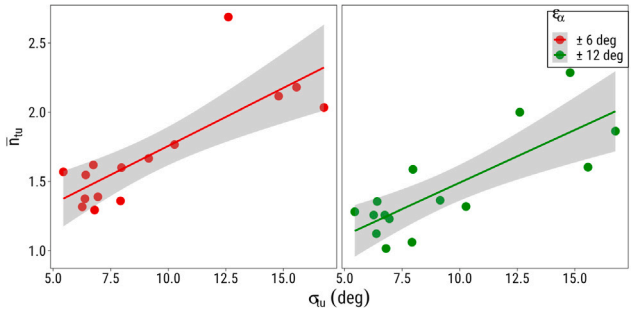

Fig. C.7. Average number of key presses required to complete a trial in the Turn action task as a function of turn variability for trials of tolerances of ±6 and ±12 degrees. Each point corresponds to a subject. The line represents the best fitting linear model.

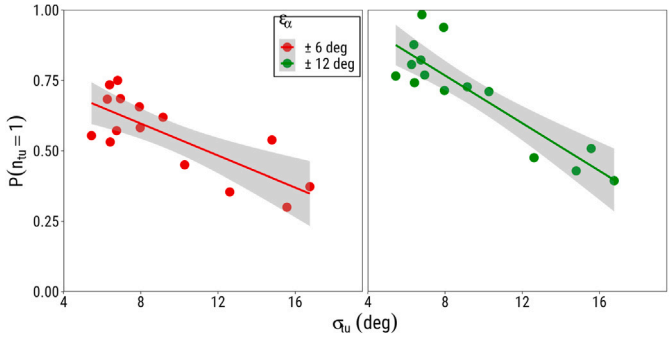

Fig. C.8. Relative frequency of completing a turn in a single key press as a function of turn variability for trials of tolerances of ±6 and ±12 degrees. Each points corresponds to a subject. The linear represents the best fitting linear model.

C.4. Space Track

In Fig. C.6 we plot the asymptotic score estimate in Space Track, ordered according to subject ranking.

There were large individual differences in main task performance.

C.5. Relationship between tasks

There was a strong correlation between long turn variability σ_{tu} and the two turn action measures, average number of key presses in a turn action $\bar{n}_{tu}$ (Fig. C.7) and probability of completing the turn with a single key press $P(n_{tu} = 1)$ (Fig. C.8). This is because that probability corresponds to the portion of the single turn distribution that falls within the tolerance region (see Fig. C.5 for an example of that portion for three subjects spanning the full range of σ_{tu} and both tolerance regions used in the study). Both $\bar{\imath}$ and σ_{tu} both strongly correlated with Space Track score (Fig. C.9).

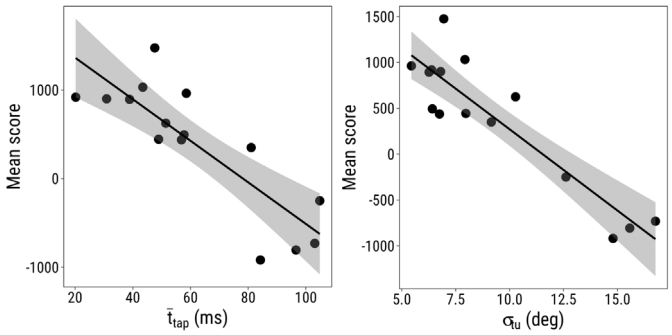

Fig. C.9. Relationship between asymptotic subject performance in Space Track (i.e., average score over best scoring 4 games) and individual difference measures. (a) Long turn variability and score. (b) Mean tap duration and score.

Appendix D. Action models

D.1. Turn action

We tested three models of how subjects orient the ship to within an aim accuracy. All models assume that the first key press is executed with the variability of a long key press σ_{tu} as measured in the Long turn task. The probability of reaching the aim accuracy with a single key press of such a variability was defined in the main text as $\hat{P}(n_{tu} = 1; \sigma_{tu}, \epsilon_\alpha)$. Here, we will use the notation P_{long} as a shorthand.

The models differ in the variability that they assume corrective turns are executed with. Some assume σ_{tu}, whereas others assume that corrective turns are brief taps and, hence, are less variable. Models are tested based on their predictions of n_{tu} for both aim tolerances of 6 and 12 degrees.

D.1.1. Model 1: All key presses are long

Model 1 assumes that each turn is equally variable with a SD of σ_{tu}. The probability of completing the turn in n_{tu} key presses equal the probability of failing the turn in the first $n_{tu} - 1$ turns and succeeding on the following one:

$$
\begin{aligned}
P_{n_{tu=1}} &= P_{long} \\
P_{n_{tu=2}} &= (1 - P_{long}) \cdot P_{long} \\
P_{n_{tu=3}} &= (1 - P_{long})^2 \cdot P_{long} \\
&\vdots \\
P_{n_{tu=i}} &= (1 - P_{long})^{(i-1)} \cdot P_{long}
\end{aligned}
$$

This results in a geometric series, which has a well-known closed-form solution, in this case for the expected number of key presses $\bar{n}_{tu}$:

$$\bar{n}_{tu} = \sum_{i=1}^{\infty} i \cdot P_{long} \cdot (1 - P_{long})^{(i-1)} = \frac{1}{P_{long}} \tag{D.1}$$

This equation predicts well the number of key presses for high ϵ_α, but overestimates it for the low one (Fig. D.1).

D.1.2. Model 2: Corrective key presses are taps

According to Model 2, all corrective turns are taps with a SD of σ_{tap} (scaled by the turn rate $\omega = 3$ deg/tick), which have a probability of succeeding P_{tap}:

$$P_{tap} = P(n_{tu} = 1; \sigma_{tap}, \epsilon_\alpha) = \int_{-\epsilon_\alpha}^{\epsilon_\alpha} \mathcal{N}(0,\ \sigma_{tap}^2) d\alpha \tag{D.2}$$

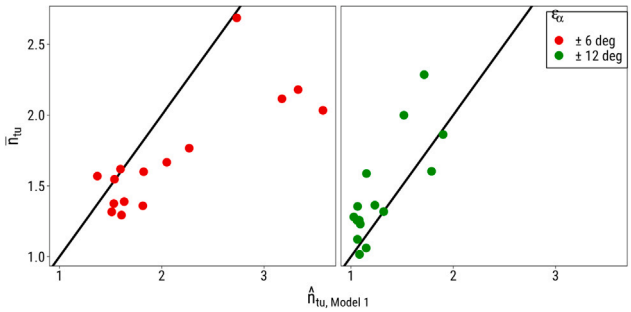

Fig. D.1. Turn action model 1: Predicting number of key presses required to complete Turn action task. Black line plots line of the equivalence.

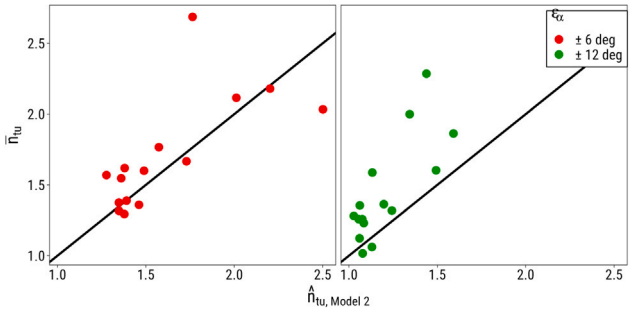

Fig. D.2. Turn action model 2: Predicting number of key presses required to complete the Turn action task. Black line plots the line of equivalence.

The probability of failing to turn to the target region on the first turn is $(1 - P_{long})$, whereas it is $(1 - P_{tap})$ on all subsequent turns:

$$
\begin{aligned}
P_{n_{tu=1}} &= P_{long} \\
P_{n_{tu=2}} &= (1 - P_{long}) \cdot P_{tap} \\
P_{n_{tu=3}} &= (1 - P_{long})(1 - P_{tap}) \cdot P_{tap} \\
&\vdots \\
P_{n_{tu=i}} &= (1 - P_{long}) \cdot (1 - P_{tap})^{(i-2)} \cdot P_{tap}
\end{aligned}
$$

The expected number of key presses $\bar{n}_{tu}$ can be rearranged to include a geometric series, which has a closed-form solution too:

$$
\begin{aligned}
\bar{n}_{tu} &= P_{long} + (1 - P_{long}) \left(\sum_{i=2}^{\infty} i(1 - P_{tap})^{i-2} P_{tap} \right) \\
&= P + \frac{1 - P_{long}}{1 - P_{tap}} \left(\frac{1}{P_{tap}} - P_{tap} \right)
\end{aligned} \tag{D.3}
$$

Model 2 predicts well the number of key presses for the more narrow tolerance region, but underestimates them for the wider one (Fig. D.2).

D.1.3. Model 3: Corrective key presses are either taps or long key presses

Model 3 assumes that corrective turns can be either taps or long key presses. However, it does not assume that the choice of which to use is random. Instead, it assumes that whichever key press brings the ship closer to the target orientation will be chosen.

Which key press will bring the ship closer is based on the angle change each key press leads to. For taps, this is assumed to be $\bar{t}_{tap}$, scaled by the turn rate ω. For long key presses, this is assumed to be the time difference between the moment the "press" command is sent and the moment the "release" command is sent according to ACT-R (i.e., 130 ms). The final equation is the same as for Model 2, with P_{tap} modified to P_{corr}:

$$P_{corr} = \left(P_{tap} \cdot \int_{\epsilon_\alpha}^{mid\,point} \mathcal{N}(0,\ \sigma_{tu}^2) d\alpha + P_{long} \cdot \int_{mid\,point}^{\infty} \mathcal{N}(0,\ \sigma_{tu}^2) d\alpha \right)$$

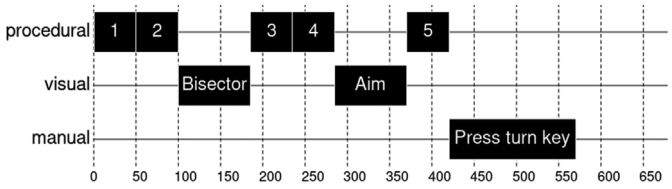

Fig. E.1. Process trace of ACT-R model of first turn. The x-axis plots time (in ms), whereas the y-axis represents the activity of the various modules of ACT-R that are involved in this model. Before turning the model needs to estimate the ship target orientation (i.e., the bisector), perceive the aim and then press the appropriate turn key. Productions 1 and 2 direct attention to the inter-segment region and command the visual module to perceive the region where the two segments intersect. This is followed by the visual module attending to that region. Following that, production 3 fires, which extrapolates the bisector and stores it as the target orientation value. In addition, this production directs visual attention to ship aim. Then, production 4 commands the visual module to encode the current ship aim, followed by the actual encoding. Finally, production 5 compares the target ship aim and the current ship aim and executes the appropriate key press.

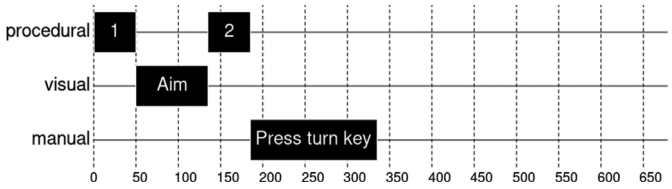

Fig. E.2. Process trace of ACT-R model of turn corrections. The x-axis plots time (in ms), whereas the y-axis represents the activity of the various modules of ACT-R that are involved in this model. The model starts with a target aim already stored. Additionally, the model is already attending to ship aim. Hence, the model needs to only encode the current ship aim value (production 1 and visual module activity) before deciding which key to press to execute the turn correction (production 2 and manual module activity).

$$/\int_{\epsilon_\alpha}^{\infty} \mathcal{N}(0,\ \sigma_{\text{tu}}^2)d\alpha \tag{D.4}$$

In this equation, the midpoint is between the short tap and 130 ms. This model has the following effect: At low ϵ_α, the ship can be closer to the target orientation while still being outside of the tolerance region. At high ϵ_α, the ship is always relatively far from the target orientation (i.e., at least 12 deg away). Consequently, a short key press is more likely to suffice at narrow tolerances than at wide ones. Its predictions are in the main text.

Appendix E. Cognitive delays of turn and thrust key presses

We estimated the cognitive delays associated with preparing each action using the ACT-R cognitive architecture. There are three cognitive costs to estimate, for which we constructed separate models. First are the costs associated with preparing the first turn (process trace: Fig. E.1). Second are the costs associated with preparing turn corrections (Fig. E.2). Last are the costs associated with preparing a thrust, whether a first thrust or a thrust correction (Fig. E.3). The reason why the first turn takes longer to prepare than turn corrections is because the model needs to first estimate its target orientation before starting to turn. This target orientation is already estimated when turn corrections need to be executed, which makes preparing turn corrections faster.

Appendix F. Velocity after thrusting

If ship orientation is ideal, the ship is aimed at the angle bisecting the inter-segment angle. The ship enters each intersection at a certain speed and aims to maintain that speed after thrusting. Thus, it thrusts for as long as it is necessary to be flying exactly along the second

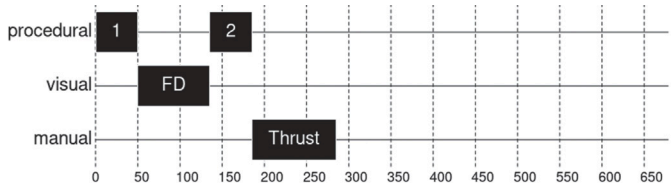

Fig. E.3. Process trace of ACT-R model of thrusts. The x-axis plots time (in ms), whereas the y-axis represents the activity of the various modules of ACT-R that are involved in this model. The model starts with a target flight direction already stored. Additionally, the model is already attending to ship flight direction. The model encodes ship flight direction (production 1 and visual module activity) and then sends a command to thrust (production 2 and manual module activity). Note that unlike turning, there is only one button that can be used for thrusting. We therefore assume that the thrust is prepared in advance and hence motor preparation time is reduced by 50 ms relative to turn motor preparation time.

segment, on average. Any offsets from that ideal trajectory are due to thrust noise. The distribution of the resultant velocities is given by:

$$\begin{aligned} v_{\text{final, x}} &\sim \mathcal{N}\left(|\vec{v}_{\text{target}}|\cos(\beta), \cos\left(\frac{180+\beta}{2}\right)\cdot a\cdot\sigma_{th\ act}\right) \\ v_{\text{final, y}} &\sim \mathcal{N}\left(|\vec{v}_{\text{target}}|\sin(\beta), \sin\left(\frac{180+\beta}{2}\right)\cdot a\cdot\sigma_{th\ act}\right), \end{aligned} \tag{F.1}$$

where the a is the acceleration value and $\frac{180+\beta}{2}$ is the orientation of the bisector. The ship's flight direction distribution can be computed from the velocity distribution, whereby each velocity value is transformed to a flight direction value ($\arctan(v_x/v_y)$). When its aim is not ideal, the ship is not thrusting exactly at $\frac{180+\beta}{2}$ and the resultant velocity distribution is adjusted accordingly. In that case, the ship needs to also thrust for a little less or more than it would ideally. The general solution to thrust time is given by:

$$t_{\text{target}} = \frac{v(\cos(\alpha)-1)}{a\cdot\cos(90+\alpha/2+\Delta\alpha)} \tag{F.2}$$

Note that the analysis above assumes that thrust duration is a real value, whereas it is discrete in the game, limited by game frame rate. The actual simulations take into account only discrete thrust values.

Appendix G. Study of the robustness of model predictions

In the main text, we evaluate model predictions relative to asymptotic performance estimated as the 4 consecutive games with the highest score. Mean score and number of crashes are estimated over this window. Here we systematically vary this window size between 1 and 10 games to demonstrate that the results in the main text are robust. Overall, the correlation between model prediction and data remain stable, as do the MAE and MSD values. Moreover, as indicated by MSD, the model performs higher than subjects, the more games we include. This is likely a consequence of including games, in which subjects are not very skilled yet (see Table G.1).

Appendix H. Relationship between measures when outlier subjects are included

In this section we provide justification that the two outlier subjects responded unrealistically slow in the Tap task. As indicated in the main text, their mean tap times were at least twice as slow as those of any other subject (see x-axis of Fig. H.1). Moreover, in absolute terms, the tap durations are excessively high. Moreover, they are not commensurate with subject performance in other tasks. For example, as indicated in Table 2, there is a highly significant positive correlation between long turn variability and mean tap duration (see green points in Fig. H.1). When the outlier subjects are included, the correlation drops to $r(\bar{t}_{tap}, \sigma_\alpha) = .38$ and is no longer significant. The same change is observed for the correlations between the turn action task measures and

Table G.1
Study of the robustness of resource-rational model evaluation as a function of the window size (in terms of number of consecutive games) over which Space Track asymptotic performance is computed. The two values compared are the asymptotic score and number of crashes predicted for each subject by the model, and the asymptotic values estimated over the said window size from data. We compared these values with three measures: the correlation r, the mean absolute error (MAE), and the mean signed deviation (MSD).

Averaging window size	1	2	3	4	5	6	7	8	9	10
r(score)	.85	.87	.87	.87	.87	.87	.87	.87	.88	.88
r(crashes)	.89	.89	.91	.86	.89	.89	.89	.88	.88	.87
MAE(score)	328	332	319	332	323	327	337	341	347	356
MAE(crashes)	2.92	2.63	2.19	1.89	2.04	2.06	1.95	2.05	1.88	1.95
MSD(score)	2	92	135	163	170	185	201	212	232	293
MSD(crashes)	.75	.25	−.03	.15	−.11	−.25	−.33	.54	−.61	−.61

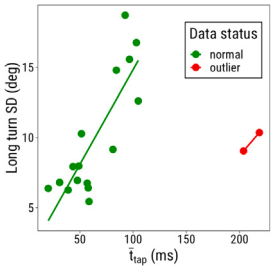

Fig. H.1. Mean tap duration and long turn standard deviation for outlier and non-outlier subjects. Each observation corresponds to a single subject. The outlier subjects demonstrate a mean tap duration much higher than the group of non-outlier subjects.

tap duration, which now drop to $r(\bar{t}_{tap}, \bar{n}_{tu,\varepsilon_\alpha=6}) = .46$, $r(\bar{t}_{tap}, \bar{n}_{tu,\varepsilon_\alpha=12}) = .30$, $r(\bar{t}_{tap}, P_{1,\varepsilon_\alpha=6}) = -.17$, $r(\bar{t}_{tap}, P_{1,\varepsilon_\alpha=12}) = -.32$ and are all no longer significant. At the same time, all correlations between long turn and turn action measures remain strong ($r(\sigma_\alpha, \bar{n}_{tu,\varepsilon_\alpha=6}) = .77$, $r(\sigma_\alpha, \bar{n}_{tu,\varepsilon_\alpha=12}) = .84$, $r(\sigma_\alpha, P_{1,\varepsilon_\alpha=6}) = -.82$, $r(\sigma_\alpha, P_{1,\varepsilon_\alpha=12}) = -.89$) and highly significant even when the outlier subjects are included.

References

Ackerman, P. L. (1987). Individual differences in skill learning: An integration of psychometric and information processing perspectives.. *Psychological Bulletin*, *102*(1), 3–27. http://dx.doi.org/10.1037/0033-2909.102.1.3.

Anderson, J. R. (1991). The adaptive nature of human categorization.. *Psychological Review*, *98*(3), 409–429. http://dx.doi.org/10.1037/0033-295x.98.3.409.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?*. Oxford University Press, http://dx.doi.org/10.1093/acprof:oso/9780195324259.001.0001.

Anderson, J. R., Betts, S., Bothell, D., Hope, R., & Lebiere, C. (2019). Learning rapid and precise skills.. *Psychological Review*, *126*(5), 727–760. http://dx.doi.org/10.1037/rev0000152.

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective.. *Psychological Review*, *96*(4), 703–719. http://dx.doi.org/10.1037/0033-295x.96.4.703.

Dimov, C. M., Khader, P. H., Marewski, J. N., & Pachur, T. (2019). How to model the neurocognitive dynamics of decision making: A methodological primer with ACT-R. *Behavior Research Methods*, *52*(2), 857–880. http://dx.doi.org/10.3758/s13428-019-01286-2.

Donchin, E. (1995). Video games as research tools: The space fortress game. *Behavior Research Methods, Instruments, and Computers*, *27*(2), 217–223. http://dx.doi.org/10.3758/bf03204735.

Dubey, R., & Griffiths, T. L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity.. *Psychological Review*, *127*(3), 455–476. http://dx.doi.org/10.1037/rev0000175.

Fu, W.-T., & Gray, W. D. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology*, *52*(3), 195–242. http://dx.doi.org/10.1016/j.cogpsych.2005.08.002.

Hoekstra, C., Martens, S., & Taatgen, N. A. (2020). A skill-based approach to modeling the attentional blink. *Topics in Cognitive Science*, *12*(3), 1030–1045. http://dx.doi.org/10.1111/tops.12514.

Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action.. *Psychological Review*, *116*(4), 717–751. http://dx.doi.org/10.1037/a0017187.

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, http://dx.doi.org/10.1017/s0140525x1900061x.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part I. Basic mechanisms.. *Psychological Review*, *104*(1), 3–65. http://dx.doi.org/10.1037/0033-295x.104.1.3.

Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing*. Academic Press, URL http://shelf2.library.cmu.edu/Tech/240474311.pdf.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection.. *Psychological Review*, *101*(4), 608–631. http://dx.doi.org/10.1037/0033-295x.101.4.608.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition.. *Psychological Review*, *109*(3), 472–491. http://dx.doi.org/10.1037/0033-295x.109.3.472.

Seow, R. Y. T., Betts, S. A., & Anderson, J. R. (2021). A decay-based account of learning and adaptation in complex skills.. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *47*(11), 1761–1791. http://dx.doi.org/10.1037/xlm0001071.

Sims, C. R., Neth, H., Jacobs, R. A., & Gray, W. D. (2013). Melioration as rational choice: Sequential decision making in uncertain environments.. *Psychological Review*, *120*(1), 139–154. http://dx.doi.org/10.1037/a0030850.