Contents lists available at ScienceDirect

Neuropsychologia

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



The sequential structure of brain activation predicts skill



John R. Anderson^{a,*}, Daniel Bothell^a, Jon M. Fincham^a, Jungaa Moon^a

^a Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, United States

ARTICLE INFO

Article history: Received 23 September 2015 Received in revised form 15 December 2015 Accepted 16 December 2015 Available online 18 December 2015

Keywords: Computational Modeling fMRI Expertize Learning

ABSTRACT

In an fMRI study, participants were trained to play a complex video game. They were scanned early and then again after substantial practice. While better players showed greater activation in one region (right dorsal striatum) their relative skill was better diagnosed by considering the sequential structure of whole brain activation. Using a cognitive model that played this game, we extracted a characterization of the mental states that are involved in playing a game and the statistical structure of the transitions among these states. There was a strong correspondence between this measure of sequential structure and the skill of different players. Using multi-voxel pattern analysis, it was possible to recognize, with relatively high accuracy, the cognitive states participants were in during particular scans. We used the sequential structure of these activation-recognized states to predict the skill of individual players. These findings indicate that important features about information-processing strategies can be identified from a model-based analysis of the sequential structure of brain activation.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

This paper is concerned with identifying the patterns of brain activation that are associated with skilled performance in a complex task. Many studies have investigated perceptual expertize in recognizing visual patterns (e.g., Bilalić et al., 2011, Scherf et al., 2011; Tarr and Gauthier, 2000). The typical result is greater activation in visual areas like the fusiform for objects in the domain of the expertize. There is also increased distinctiveness in the activation patterns in these regions (Haxby et al., 2000, Nestor et al., 2011). However, increased activation is not always a sign of increased skill. Looking at other domains such as problem solving or visual-motor tasks, researchers have found that higher skill was characterized by decreased activation, which is taken as a sign of improved processing efficiency (Anderson, 2005; Gobel et al., 2011; Poldrack et al., 1998). Chein et al. (2005, 2012) emphasize the drop out of different brain regions as a skill develops.

While activation in specific regions may reflect expertize in some tasks and patterns of activation across regions may reflect expertize in other tasks, such static perspectives ignore the importance of the sequential structure of many tasks. Behavioral studies have shown that experts often approach a task with a very different sequence of mental steps than someone less expert (Anderson et al., 1984; Larkin, 1981). This paper will describe such

E-mail address: ja@cmu.edu (J.R. Anderson).

http://dx.doi.org/10.1016/j.neuropsychologia.2015.12.014 0028-3932/© 2015 Elsevier Ltd. All rights reserved. a task where expertize is reflected dynamically in how whole brain activity changes over an episode of doing the task. That is, expertize is reflected by patterns in the spatial-temporal structure of brain activation. We will describe how these patterns can be measured by combining of multi-voxel pattern analysis (MVPA – e.g., Norman et al., 2006; Pereira et al., 2009) and hidden Markov models (HMMs, e.g., Rabiner, 1989) to recognize patterns defined in reference to a cognitive model of how the task should be performed. While the current task is a video game whose sequential structure is relatively easy to identify, the general approach is applicable to a wide range of tasks.

1.1. The Space Fortress game

The Space Fortress video game, which has had a long history in the study of skill acquisition, was first used in the 1980's by a wide consortium of researchers (e.g., Donchin, 1989; Frederiksen and White's, 1989; Gopher et al., 1989). Fig. 1 shows a schematic of the Space Fortress video game. The player is supposed to keep a ship (red) flying between the two hexagons. This is challenging because the ship is flying in a frictionless space and there is no braking system for navigation. To navigate the player must combine thrusts in various directions to achieve a circular orbit. While doing so, they need to shoot at the central fortress (blue) and avoid being shot by it. In addition the player must deal with mines (green) that appear at irregular intervals. Players must distinguish between "foe" and "friendly" mines, which can be identified by a set of letters that appear under header IFF ("Is Foe or Friendly") in



 $[\]ast$ Correspondence to: Department of Psychology \sim Baker Hall 345D, Pittsburgh, PA 15213, United States.



Fig. 1. Left: a schematic of the screen for Space Fortress. The ship (indicated in red) is controlled by the user and must circle the fortress (in blue) staying within the two hexagons. The hexagon (in green) represents a mine coming at the ship. At the bottom of the screen is an information bar that displays critical information. The red highlighted vulnerability information keeps track of the number of times the fortress has been hit and the highlighted IFF information contains a letter that identifies whether the mine is friendly or foe. The dollar sign (yellow) is one of a sequence of symbols that must be tracked for bonus points. Right: mapping of fingers onto actions – L=counter-clockwise (left), T= thrust, R=clockwise (right), F=fire, I=IFF; B=bonus. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the bottom panel (before each game the player learns 3 letters that will indicate foe mines for that game). Finally, potential bonus symbols (yellow) appear on the screen below the fortress. If players identify when two dollar signs appear consecutively they will receive points or ammunition.

The player interacts with the ship by pressing 6 keys with their ring, middle, and index fingers on both hands¹. The left hand uses the three keys involved in ship control – for turning clockwise, counterclockwise, and thrusting. The left hand needs to master the correct execution of thrusts and turns to maintain a circular orbit in the frictionless space. The right hand controls the firing and response to critical events. Timing of right hand keys is particularly important. Shots at the fortress must satisfy a number of timing constraints – at least 250 ms between shots up until the fortress is vulnerable for a "kill shot" which requires a pair of shots under 250 ms. Neutralization of foe mines requires a double click with another right-hand finger and must be between 250 ms and 400 ms. The players are trying to optimize their point totals, which is a sum of points for navigating, dealing with the fortress, processing the mines, and responding to the bonus opportunities.

A number of neuro-imaging studies have looked at learning in Space Fortress. Erickson et al. (2010) found that striatal volume predicted learning. Working with basically the same data set, Vo et al. (2011) found that both mean activation and multi-voxel patterns in the striatum predicted learning gain. Anderson et al. (2011) found that the increase in motor and fusiform activation in response to mines predicted learning gains. While we will report results on learning gains, the focus will be on differences among individuals that are stable across learning.

1.2. The cognitive model

Bothell (2010) developed a model in the ACT-R architecture (Anderson et al., 2004; Anderson, 2007) that was capable of performing the task². This model predicted the brain imaging data in Anderson et al. (2011). The model's performance is subject to human-like constraints on the speed of visual, motor, and cognitive processes. These constraints have been set on the basis of a wide literature. The model performs at a high level of performance comparable to our best human participants. However, because of its human-performance limitations, it does not do as well as a "pure-AI" system can, as we will discuss.

The ACT-R architecture consists of a set of modules, which are active in parallel to perform a complex task like this video game. Fig. 2a illustrates the modules of relevance to the model for Space Fortress and Fig. 2b illustrates how these modules are deployed in the 2 s after a mine appears. During this 2 s the Visual module is detecting the presence of a mine and encoding the letter to identify the mine; the Goal module is switching focus from processing the fortress to processing the mine; the Imaginal module is determining how to aim at the mine; the Declarative module is retrieving information about whether the letter identifies a foe: and, if appropriate, the Manual module is executing the double tap needed to enable a shot at the foe mine in addition to firing at the mine. The coordination of these modules is controlled by a set of production rules (procedural module). Each rule recognizes particular states of modules and makes requests of the modules. These productions can be broken into three categories - productions that control the navigation of the ship around the fortress, productions that fire at the fortress and collect bonus points, and productions that handle mines. All of these productions are highly reactive and respond to the state of the game. Only one production is able to apply in at any time. When multiple productions could apply, the most urgent is determined by ACT-R's conflict resolution.

1.3. The Current Study

This paper will use that model to assess how participants are responding to the game and will use this information to predict their skill. The current experiment differs from our earlier 2011 study in three ways that facilitate the goals of this paper. First, the number of participants was doubled to get a better sample of individual variation. Second, whole brain activation was used rather than focusing on activation in specific brain regions. Third, rather than having a very regular sequence of game events (which facilitated model fitting in our earlier study) an asynchronous sequence was used to better assess sequential structure.

The past study showed that it was informative to examine performance in various simpler versions of the game as well as performance in the most complex version of the game, and that

¹ Our implementation uses the Pygame version of Destefano and Gray (2008) and involves a significant control simplification over the original task, which required mastering a 2nd order, acceleration-control joystick.

² The runnable model and a detailed description of that model and its behavioral correspondence are available in the supplementary materials and at http://act-r.psy.cmu.edu/publications/pubinfo.php?id=974. Also we have placed at the web site experimental software that runs with human participants or the model.



Fig. 2. (a) The interaction of the ACT-R modules during the presentation of a mine. Beside each module is an example of its activity during this period. (b) A swimlane representation of the activity of the ACT-R modules during the 2 s after a mine appears. This would correspond to an IFF scan in the state analysis to follow.

same condition structure will be used here. In the **Orbit** condition all that participants have to do is successfully fly their ship around the fortress, receiving points for keeping the ship within the outer hexagon, not hitting the inner hexagon, and keeping their speed within the specified bounds. In the **Mines-Only** condition participants also have to deal with mines, receiving additional points for successfully destroying mines but losing points if the mines hit them. In the **Fortress-Only** condition participants have the fortress and bonus symbol, receiving additional points for destroying the fortress and collecting bonus points, but losing points if the fortress shoots them. The **Both** condition is the original Space Fortress game, involving the points for navigation, mines, destroying fortress, and bonus symbols.

2. Material and methods

2.1. Participants

Forty right-handed participants (14 female and 26 male, 18–31 years of age, mean=24.0) were recruited from the Pittsburgh community.

2.2. Experimental procedure

Fig. 3 unpacks the overall structure of the experiment. The experiment lasted 5 days with two 1-hour sessions each day. The second session of the first and fifth day was performed in an fMRI



Fig. 3. The structure of the experiment: within days there are 2 sessions and within each session there is a sequence of 16 games. The 4 games in each quarter of each fMRI session involved a random sequence of the Orbit, Mines, Fortress, and Both conditions (the order in the figure is just illustrative). Two 3-minute games occurred in each 8-minute scanning block. Each block had a 20-second fixation period (denoted by F) at the beginning, middle, and end. Each Game was preceded by 10 s to study the Foe letters (denoted by L) and then 10 s to inform participants about the condition (denoted by C). Each game was followed by 10 seconds of feedback on score (denoted by S).

Experiment

magnet.

The Space Fortress game was the same as in Anderson et al. (2011) with one difference that allowed us to better assess sequential structure: rather than mines appearing regularly once every 20 s they appeared at random intervals uniformly sampled from 4–8 s. In all sessions participants responded with three fingers (ring, middle, and index) in two data gloves (see Fig. 1 for mapping of fingers). In each session they performed 8 two-game scanning blocks for a total of 16 games per session. Each game began with 20 s of fixation, then 10 s to memorize the 3 letters that would be used to identify "foe" mines in that game, and then 10 s notifying them of the condition (Orbit, Mines-Only, Fortress-Only, or Both) of the upcoming game. The game lasted for 3 min and was followed by 10 s in which participants could study their score. A final 14-second fixation period followed the feedback for the second game in a block.

There were 10 sessions, administered 2 per day. On the first and fifth day, there were 4 games in each of the Orbit, Mines-Only, Fortress-Only, and Both conditions. Each game appeared once in random order in each set of four games in a session. The second sessions of the first and fifth days (Sessions between 2 and 10) were conducted in the scanner. Half of the participants practiced the Both condition during the intervening 3 days and half alternated between the Mines-Only and Fortress-Only conditions during those days.

2.3. Imaging analysis

Images were acquired using gradient echo–echo planar image (EPI) acquisition on a 3 T Verio Scanner using a 32-channel RF head coil, with 2 s. repetition time (TR), 30 ms echo time (TE), 79 degree flip angle, and 20 cm field of view (FOV). We acquired 34 axial slices on each TR using a 3.2 mm thick, 64×64 matrix. This produces voxels that are 3.2 mm high and 3.125×3.125 mm². The anterior commissure-posterior commissure (AC-PC) line was on the 11th slice from the bottom scan slice. Acquired images were pre-processed and analyzed using AFNI (Cox, 1996; Cox and Hyde, 1997). Functional images were motion-corrected using 6-parameter 3D registration. All images were then slice-time centered at 1 s and co-registered to a common reference structural MRI by means of a 12-parameter 3D registration and smoothed with a 6 mm full-width-at-half-maximum 3D Gaussian filter to accommodate individual differences in anatomy.

We calculated the mean maximum movement per block. Roll, pitch and yaw were (in degrees) 0.66, 1.65, and 0.80, respectively. Translations in *z*, *x* and *y* directions were (in mm) 2.06, 0.77 and 0.94 respectively. Only yaw showed a significant decrease from Session 2 to Session 10 (from 90 degrees to 71 degrees; t(35) = 2.24, p < 0.05) and it was not significantly correlated with score.

In complex tasks like this, which involve many cognitive processes, the focus is on patterns of activity across the whole brain common across individuals. As a step of dimension reduction and to accommodate variations in anatomy over participants that may not be dealt with in co-registration, we aggregated the voxels in each slice into larger 2×2 voxel regions. The fundamental motivation is to predict across individuals and this eliminates our ability to have very precise localization. The process of co-registration in itself already forgoes strong inference on smaller subcortical structures, independent of other dimension reductions that we may apply. We chose only to aggregate in the 2D plane because the mismatch between subject brains can be dramatic at the curvature around inferior frontal brain regions. In this area, moving from one slice to the next in the inferior direction, there is often an abrupt transition from signal to non-signal (imagine the saggital view). There are 12,481 such 2×2 regions. However, some of these show an excess of extreme values for some participants, likely reflecting differences in anatomy³. Most of these outlying regions were located on the top and bottom slices as well as near the perimeter of the brain. Eliminating these resulted in 9,029 usable regions for analyses.

To correct for brain-wide activation fluctuations that were not task related (for instance, due to changes in breathing patterns found while playing video games, Birn et al., 2008), we regressed out correlated activity involving bilateral regions of the secondary auditory cortex (two bilateral rectangular volumes defined as $5\times5\times4$ voxels in the functional acquisition matrix centered at Talairach coordinates \pm 42, -61, -9). These regions correspond to the Aural module in ACT-R and should not be engaged in the Space Fortress task. We regressed the full 8-minute BOLD response in each of the 9029 regions against the mean BOLD response in these reference regions. The corrected activity for a region of interest was the residual BOLD response not predicted plus the mean BOLD signal in that region for that 8-minute period. Anderson et al. (2011) found that this residual measure did not change the direction of any of the effects but increased the significance of many. Performing this step of analysis keeps our processing with the earlier publication⁴.

The means of the scans 4 through 6 in each fixation period were used to establish baselines for each game. The baseline for any scan in the first game of a block was calculated as the linear trend between the first and second fixation means. The baseline for the second game was calculated as the linear trend between the second and third means. The BOLD response on any game scan was calculated as its percent change from its baseline. A Wiener filter (Glover, 1999) with a noise parameter of .1 was used to deconvolve the BOLD response into an inferred activity signal on each scan. The hemodynamic function used in this deconvolution is the SPM difference of gammas (Friston et al., 2011: gamma(6,1)gamma(16,1)/6). The Appendix reports an exploration of the effects of variations on this deconvolution step on classification and provides a justification of the hemodynamic function we used.

The output of this process is an estimation of activity in 9029 regions during each of the 90 two-second scans in a game. However, because of high correlation among regions this reflects nothing like 9029 independent pieces of information. For the purpose of identifying sequential patterns, we performed a spatial principal component analysis (PCA) of the voxel activity where each voxel was treated as a variable that varies over scans, games, and participants. Separate PCAs were done for Sessions between 2 and 10. The first 100 components of these PCAs were used in subsequent analyses. These captured just over 73% of the total variance in each session. To focus on the sequential changes across the course of games (rather than differences among games within a condition), the 100 PCA scores for each game were standardized to have means equal to the mean for that condition and variance 1. The predictive analyses will use these normalized PCA scores over the 90 2-second scans that define a game.

3. Results

3.1. Preliminary analyses

Our focus will be on relating participants' average scores to

 $^{^3}$ The definition of a "excess of extreme values" is more than .5% of the scans with values more than 10% from baseline.

⁴ To review the considerations that drove this choice, the problem is that some participants will change their breathing patterns when the game playing gets challenging and this has identifiable effects on the BOLD response. We did try a number of alternative regions but found that using a grey-matter are, that did not respond to task structure, like the auditory region, provided the best correlate.

their brain activation⁵. Fig. 4 displays the average scores over 16 games for each participant in the two imaging sessions, Session 2 and Session 10. There is substantial improvement over the course of the experiment (average 2494 points in Session 2 and 3518 points in Session 10). All but one participant improved and the correlation between participants' score on scanner session 2 and their score on scanner session 10 was 0.752. Of the 20 top players in Session 2, 17 are still in the top 20 on Session 10. Fig. 4 displays the individual scores on these two days plus the results from 20 simulated participants (ACT-R model), which averaged 4190 points. The simulated scores are more tightly clustered than the human participants, but there is some variation in the model performance depending largely on randomness in the appearance of mines and bonus points and on randomness in ACT-R's motor timing. The best performing simulated participant slightly exceeds the best performing human participant (4370 points versus 4258 points). We determined what an Ideal model would do that embodied the same strategy as ACT-R but had no randomness or delays. Thus, it performed every action perfectly and could respond instantly to any event. Such a model would have averaged 5785 points given the statistics of the games that the ACT-R model faced.

The imaging analyses excluded 4 participants who showed an excess of extreme values (defined as BOLD responses greater than 5% from baseline) and thus involved 36 participants. The Appendix reports an exploratory analysis looking for regions that significantly (brain-wise alpha of 0.01) differed according to four binary variables – whether it was Session 2 or Session 10, whether the fortress was present or not, whether the mine was present or not, and whether or not the participant was in the top half of the players. Numerous sensible regions were identified for the first three binary contrasts but no region was found that separated stronger from weaker players.

Having failed to find any region that responded to the binary contrast of stronger versus weaker players, we performed a regression analysis examining the relationship between player mean activations and player average score. This analysis yielded one region in the right dorsal striatum (details of analysis in Appendix), involving the caudate and claustrum whose mean activity did correlate with score. This region appears to partially overlap with the region reported in Ref. Vo et al. (2011).

3.2. Using sequential structure of brain activation to predict expertize

Below we describe the three steps in producing a predictive measure of the sequential structure of brain activation during a game. First, independent of brain activation, we identify the sequential structure of the states occurring during game play and show that better participants display sequences of states more like the sequences that the ACT-R model displays. Second, we show that we can use brain activation to predict the state that a participant is in on any particular scan. Third, we combine the previous two steps to use the sequential structure of brain activation to diagnose the skill of a player.

3.2.1 Step 1: using a states to measure sequential structure⁶ The sequential structure of the model is controlled by what



Fig. 4. Performance of the 40 participants during Sessions between 2 and 10 plus 20 simulated participants.

production rules match the state of the game, which is recorded at any point in time. Thus, we can infer from the state of a player's game what rules would be firing if the participant were identical to the ACT-R model. As described earlier, the ACT-R productions break into three groups: navigating, handling events associated with the presence of a fortress and bonus points, and handling events associated with mines. Each of these task groups was partitioned into a unique set of states where different subsets of production rules apply (see Table 1). In the discussion to follow we will use states to refer to the four partitions of each task group. For example, for the mine-handling group the four states are: i) IFF, which refers to executing a timed double-click to neutralize a mine identified as foe so that it can to be shot; ii) **Shoot** the mine; iii) **Fix the orbit** from drift during processing the mine; and iv) None when no mine is present. The same states for each group will occur for all players, whether they are behaving like ACT-R or not. However, their duration and sequence will be strongly influenced by the actions of the player⁷. Any particular point during a Both game can be described as a 3-tuple whose elements consist of a state from each of the task groups Navigation, Fortress and Mines. For example, a scan flying at normal speed, with a bonus opportunity, and a mine to shoot would be described by the 3-tuple < Normal, Bonus, Shoot > . We labeled each 2-second scan (for model or participant) in this manner. Because game states can change at any time, more than one state within a group can occur during a single scan. In such cases we labeled that scan with the more urgent state (as determined by ACT-R's conflict resolution, with urgency decreasing from states a to d in Table 1).

The behavior of the ACT-R model defines a semi-Markov structure (SMS – Yu, 2010) on the states in each task group. Fig. 5 shows the structure for the four Mine states, with part (a) showing the transition probabilities and part (b) showing the durations in each state (the Appendix goes into the game details to explain these probabilities and durations). These transitions and durations are determined by how a participant (human or model) plays the game. Fig. 5a contrasts the transition probabilities for ACT-R with the transition probabilities of the Ideal model (in parenthesis). Not constrained by human limitations, the Ideal model responds instantaneously and always performs actions with perfect precision.

To measure the degree to which a human's game was being played like the ACT-R model would play, we applied the ACT-R SMSs⁸ to calculate the probability that the sequence of states in that game came from ACT-R. Consider the following starting sequence of 5 Mine states: 2 scans in the None state, 2 in the Shoot state, and 1 in the None state,

The probability of this sequence is the probability that the ACT-

⁵ We are collapsing our results over the two groups. There was neither a significant difference between groups in their average scores (F(1,39)=2.31, p > 0.1) nor any significant interactions of group with the other factors, and results are averaged across group in all subsequent analyses.

⁶ The analyses in the remainder of the section is unconventional. The underlying data, software and code is available and discussed at http://act-r.psy.cmu.edu/ ?post_type=publications&p=16893.

⁷ The game itself also strongly influences the duration and sequence of these states.

⁸ The Navigation SMS applies over all conditions; the Mine SMS only applies in the Mine and Both conditions; the Fortress SMS only applies in the Fortress and Both conditions. The likelihood of a game is the average of the SMSs that apply.

Table 1

Three groups of game states: descriptions of game situations that evoke different productions in each state

1. Navigation.

Start. When the game begins or the ship respawns after being destroyed, the ship starts out motionless and needs to be put into orbit around the fortress.Fast. When the ship is in orbit with a speed above 1.7 pixels per second.Slow. When the ship is in orbit with a speed below 1.0 pixels per second.Normal. When the ship is in orbit with a speed between 1.0 and 1.7 pixels per second.

2. Fortress

- Bonus. The bonus symbol repeats and there is an opportunity to collect bonus points.
- Kill. The fortress has been hit 9 or more times, the ship is aiming at the fortress, and the special hit shot can be executed.
- Shoot. The fortress has been fewer than 9 times, the ship is aiming at the fortress, the fortress and can be shot.
- None. The ship is not aiming at the fortress and there is not a bonus opportunity.

3. Mines

- **IFF.** An enemy mine is present which requires a double-click to be neutralized so that it can be shot.
- **Shoot.** A friendly mine is present or an IFF'd mine is present, both of which can be shot.

Fix. After destroying a mine the ship has drifted out of its circular orbit. **None.** No mine is present and the ship has not drifted out of its circular orbit.

R model stays in the None state for 2 scans (0.48) times the probability of it transitioning to the Shoot state (0.5), times the probability of it staying in the state for two scans (0.46), times the probability of it going to the None state (0.65). This leads to a probability of 0.072 for that sequence. The probabilities get very small when one calculates sequences of 90 states in a game and so we calculate log probabilities. This also proves to be a better scale for correlating with score.

The average log probability of the games from a player was used as the measure of similarity of that player to ACT-R. Fig. 6a shows how the average score for different participants varies with the average log probability, plotting separately participants for Session between 2 and 10 as well as the 20 simulated participants. The correlation between log probability and score for human participants is 0.899 (the within-session correlations are 0.850 for Session 2 and 0.749 for Session 10). We used three permutations tests of 10,000 permutations to establish that the 0.899 was significantly greater than three null models. First, to test that the correlations were greater than a chance assignment of states to scans, we permuted all the states for all the subjects across days. The mean correlation was indeed 0 and the maximum correlation across the 10,000 permutations was 0.438. Second, to determine that the effects were not just due to changes in the states across days, we permuted the ACT-R state sequences for each day. The mean correlation in this case was 0.622 and the maximum of 10,000 was 0.746. Third, to determine that the effect was just due to the states in a game their sequence, we randomized the state sequences for each game for each participant for each day. The mean correlation was 0.844 and the maximum in 10000 was 0.887. So, in all 10,000 cases for all 3 permutations, correlations that did not reflect the sequential structure of the states were less than the observed 0.899. The relatively high correlations for within-game permutations reflects the fact that there is strong correlation between the states one visits and the sequential structure one displays.

We also looked at the correlations for other models of behavior: Log probability based on the Ideal model correlates 0.725 and log probability based on average Session 10 behavior drops to 0.678. Using tests for paired correlations (Steiger, 1980), these two predictors are not significantly different (t(70) = 1.16, p > 0.2) but both are significantly worse than the cognitive model used for Fig. 6a (difference with Ideal model is t(70) = 10.95, p < 0.0001; difference with Session 10 behavior is t(70) = 8.88, p < 0.0001). The X-axis in Fig. 6a reflects an average of the log odds of the SMS's for navigation, fortress, and mines. Individually, each has a strong correlation to score (navigation 0.747, fortress 0.799, and mines 0.812). In a stepwise regression, all three factors make independent contributions to predicting score. The optimal weight of these three factors would be 0.26*log-navigation+0.35*logfortress+0.39*log-mines, which would result in a very slightly stronger correlation with score (r=0.902) than the simple equal weighting we are using.

Part b of Fig. 6 shows the relationship between the increase in the log probability from Session 2 to Session 10 and the increase in score (r=0.704). To assess the overall statistical significance of the relationships displayed in Fig. 6s we fit a regression model to the data of the form Score=a+b*LP where LP is log probability. The best fitting parameters where a=8500 and b=56. The slope b is highly significant (t(70)=17.18, p < 0.0001). The fit is not improved by having different intercepts (a's) for the two sessions (t(69)= 1.11, p > 0.10) or by having different slopes(b's) for the two sessions (t(69)= 1.34, p > 0.10).

Thus, not only do the better participants have scores like the ACT-R model, they are achieving this performance in similar ways. Moreover, they improve their score over sessions to the degree they come to act more like the model. It is worth noting that the relationship in Fig. 6 between log probability and score emerges without any parameter estimation specific to predicting score. The model and its parameters are the same as those used in Anderson et al. (2011).



The previous section showed that we can predict the score of a



Fig. 5. (a) Transition probabilities of ACT-R model among Mine states (Ideal model probabilities in parentheses). (b) Mean number of scans the ACT-R model stays in a state upon a visit.



Fig. 6. (a) Average score of games as a function of log probability of the observed states in the games according to the ACT-R model (averaged over all SMSs that apply to a game). (b) Improvement of participants from Session 2 to Session 10 as a function of the increase in the log probability of their observed states.

participant from the sequential structure of the states they go through. We now investigate whether we can train a classifier to recognize the state a participant is in on a scan from the pattern of fMRI activity on that scan. To classify the states of specific scans we used a linear discriminant analysis (LDA), a technique we have used to classify such sequential states in other efforts (e.g., Anderson et al., 2010, 2012). As described in the methods, we used the first 100 dimensions of standardized PCA scores and the PCA was done on the scan activity deconvolved using a Weiner filter. Note that these measures were standardized such that each individual and game type had the same average values over scans. Thus, our classification effort is focused on differences among scans within a game. As in our earlier efforts, to predict the data for a particular game, we trained the LDA classifier using an equal weighting of other games (but not the test game) from that participant and sequences from other participants. This tends to give an optimal combination of sensitivity to the specifics of a particular participant and robustness based on the general predictiveness of the fMRI signal.

We trained 3 classifiers, one for each of the three task groups (Table 1: Navigation, Fortress-handling, Mine-handling), to discriminate scans among the 4 states within that group. Fig. 7 shows the success in making these classifications within each group. The average accuracy in Session 2 was 48.5% and in Session 10 it was 50.7%, which are both better than chance (25%). All participants were predicted better than chance for all groups of states. So, the brain activation for a scan does contain information about the



Fig. 7. Average accuracy in the 4-way classification of Navigation states, Fortress states, and Mine States for Session 2 and Session 10. The error bars plot \pm 2 standard errors of the mean.

state of the game at that time. While this information is far from perfect as indicated by the roughly 50% accuracy, it is sufficient to lead to success in the next step where we try to predict individual proficiency.

Our method involved training the same classifier in multiple conditions of the experiment and applying it to multiple conditions. For instance, the same Navigation classifier is used in all of the Orbit, Mines-Only, Fortress-Only, and Both conditions. This approach works slightly better than using a different classifier for each condition. The Appendix provides information on how well a classifier trained in one condition generalizes to another condition.

As an illustration of how these states are being classified, Fig. 8 shows a projection of the activation patterns for the Mine states onto the two-dimensional plane that captures the majority of the variance (79%) among the means of the four states. The Appendix provides similar plots for the Navigation States and the Fortress states. The origin of the graph corresponds to the average activation. The individual points in the space represent the average activation of different states (plotting separately between Session 2 and Session 10). The activation pattern for any point in this space can be reconstructed from the sum of three images: the origin, the x-vector weighted by the x coordinate of the point, and the y vector weighted by the y coordinate. The x-vector, on which the



Fig. 8. Projection of the 4 Mine states onto 2 dimensions. The origin of this space corresponds to the average activation during Mine games. The X and Y vectors indicate the activity added to this average activation for individuals types of scans.



Fig. 9. (a) Average score of games as a function of log probability of the LDA-inferred states in the games according to the ACT-R model (averaged over all SMSs that apply to a game). (b) Improvement of participants from Session 2 to Session 10 as a function of the increase in the log probability of their inferred states.

Shoot and IFF states have highest values, is most associated with increased left motor activation reflecting the increased right-hand actions in the states. The y-vector, on which the state of fixing the orbit is highest, has reduced left motor activity and increased posterior parietal and fusiform activity. During orbit correction there is no right hand activity and increased perceptual processing.

Successful classification of states could arise from any regularity between activation patterns and the states. However, examination the classifier solutions as shown in Fig. 8 suggests that successful classification depends on the fact that participants are exhibiting similar behavior to the model in each state. Participants who are more accurately classified tend to score higher – the correlation between a participant's mean D-prime measure (see Appendix for its calculation) of classification accuracy and score is 0.659 (t(34)=5.10; p < 0.0001) for Session 2 and 0.348 (t(34)=2.16; p < 0.05) for Session 10.

3.2.3 Step 3: diagnosing expertize from sequences of activation patterns

Now we will combine the previous two steps by classifying the brain activity on scans as to states (Step 2) and then using the sequence of classified states to predict skill (Step 1). Fig. 9a shows how the average log probability of all the games played by the participant correlates with a participant's average score across the two sessions. Fig. 9b shows how the change in log probability (from Session 2 to 10) relates to the increase in score. The overall correlation in part a is 0.798 (the within-session correlations are

0.643 for Session 2 and 0.565 for Session 10). The correlation in Part b is 0.536 between increase in score and increase in log probability. To assess the overall statistical significance of the relationships displayed in Fig. 9 we fit the same regression model as for Fig. 6: Score= $a+b^*LP$ where LP is log probability. The best fitting parameters where a=12650 and b=58. The slope b is highly significant (t(70)=11.05, p < 0.0001). The fit is not improved by having different intercepts for the two sessions (t(69)=1.45, p > 0.10) nor by have different slopes for the two sessions (t(69)=1.53, p > 0.10).

3.3. Alternatives to sequential structure

Fig. 9 shows that we can do a good job of identifying the expertize of individuals by a measure of the sequential structure of states classified from their brain activation. We considered a number of ways to identify expertize from brain activation that did not include sequential structure. Earlier we reported finding a dorsal striatum region whose mean activation correlated with score. This region was identified looking at correlation between mean activation across both sessions and mean score across both sessions. Fig. 10a shows what happens when we break this into sessions for comparison with Fig. 9a. The correlation is much reduced to just .306. There is stronger within-session correlation between this activation measure and score (0.519 for Session 2, and 0.346 for Session 10) but the mean activation does not change between sessions (0.30% on Session 2 and 0.27% on Session 10) and so it cannot explain the improvement in score.



Fig. 10. (a) Average score of games as a function of the score predicted from the participants' distribution of states. (b) Improvement of participants from Session 2 to Session 10 as a function of the increase in predicted score.

Of the various other non-sequential analyses we tried, we will describe here the only one that had some success in predicting both the within and across-session effects. This analysis used the MVPA classification of scans from Step 2. However, rather than using the sequential structure of the ACT-R SMS's from Step 1, we looked at how individual states in the ACT-R model were associated with changes to the score. Each scan resulted in an increment to the score and the total game score was just the sum of these increments. We estimated from the ACT-R model how the expected increment on a scan varied with different states. With these estimates we could then predict the score for any scan of any game. For instance, if a scan from a Both game was in the Slow Navigation State, the Kill Fortress State, and None Mine state its predicted score would be

$$m$$
 – Both + D – Slow + D – Kill + D – None

where μ -Both was the overall mean score increment for scans in the Both condition (60.9 according to the model), Δ -Fast is the mean difference between this average increment and the increment on Fast scans (-12.9 in the model), Δ -Kill is the difference associated with Kill scans (+20.1 in model), and Δ -None the difference associated with None scans (-17.7 in model). Thus the predicted value for the scan would be $60.9 \pm 12.9 + 20.1 \pm 17.7 = 50.4$. Predictions calculated in this way correlated 0.74 with the actual increments on the 28800 scans in the model games.

Having these estimates of value of game type (e.g., μ -Both) and effect of state (e.g. Δ -Kill) from the model, we could then predict the score increment for every scan of every participant game from the states on those scans classified from their brain activations. We could then sum these increments for each game to get a score for the game and average these to get a predicted score for each participant. Thus, we went from brain activation to state classification to prediction of score without considering the sequential structure of the states. Fig. 10b shows how average state-predicted score correlates with a participant's average score across the two sessions. Compared to Fig. 9, the overall correlation has dropped from 0.798 to 0.656. The reduction from 0.798b to 0.656 is significant by a test for paired correlations (Steiger, 1980) – t(70)= 2.60, p < 0.01.

To assess the overall statistical significance we fit a linear model as we had for Fig. 9a: Score=a+b*Pred where *a* is the intercept and Pred is the predicted score. There is a significant contribution of the state-predicted score (t(70)=3.88, p < 0.0005 for slope b) but also a significant contribution of having different intercepts for the two sessions (t(70)=3.99, p < 0.0005). Thus, unlike our use of the likelihood of the state sequences in Fig. 9, this measure cannot explain the entire improvement from Session 2 to Session 10.

To understand the relationship between these various approaches to predicting performance, we performed a stepwise regression in which we had included the three variables of log likelihood of the classified sequences (Fig. 9a), average activation (Fig. 10a), and average score from the classified states (Fig. 10b). Also we allowed for a different intercept for sessions. Log-Likelihood of the classified sequences entered first. It also entered with high significance if we forced all the other three variables to enter before it (t(68) = 6.15, p < 0.0001). If it entered first, the only variable that had a significant further contribution was average activation in the dorsal striatum region (t(70) = 3.76, p < 0.0005). The other two factors of session and average state-predicted score were not significant (both with nearly identical t's of 1.45, p > 0.1). Given that the MVPA was performed on PCA values normalized to have the same average values per subject, it is not surprising that mean activation before normalization still had some small

contribution to make toward predicting score.

4. Discussion

This paper has shown that an individual's skill can be predicted from how that individual's brain activation matches a model of expert performance. Stated at this level of abstraction, it may seem evident that this should be possible, but what is critical is how it was achieved. We used an ACT-R model of this task and the success of this effort provides evidence for the model and its usefulness. However, the basics of the approach could be applied with other computational frameworks if one can define a mapping between the detailed processing of such a model (e.g. Fig. 2) and a state characterization appropriate for relating to fMRI data. As we found, to capture the richness of what is happening in these scans it was necessary to use multiple parallel state classifications. Our scans could be classified by their navigational features, their mine status, and their fortress status. One can then use standard MVPA techniques to recognize these states in the brain activation. While having such a state characterization offers some predictive power, there is a critical contribution of the sequential structure of these states. To capture this sequential structure we built SMS's that embodied the model's statistics about state occupancies and transitions. The critical measure we derived was log probability - a model-based measure of the likelihood of the state sequences reflected in the brain activity. A notable feature of this approach is it never considered score in the estimation of its parameters and yet this log probability measure showed a 0.8 correlation with score.

Fig. 9b shows that the improvement of individuals can be predicted by how much more they come to behave like the ACT-R model. This is different than a number of other efforts that have tried to use early activity in regions to predict how much an individual would learn. We performed a brain-wise analysis looking for any regions that showed a significant correlation with improvement but failed to find any that showed a brain-wise alpha of .01 between activity and improvement in score. Two of the studies mentioned in the introduction (Erickson et al., 2010; Vo et al., 2011) found a significant relationship between this difference measure of improvement and regions in the striatum. Vo et al. report a positive relationship between learning and dorsal striatal activity (although they focus on the results from a support vector analysis). None of the sub-regions they considered, including the dorsal striatum, showed a significant positive correlation with the difference measure. We did find a region that overlapped with the striatum that showed a significant relationship with average score. We examined what its relationship was to improvement in score. In contrast to their result, this region in our study shows a negative correlation (r = -0.429, t(34) = 2.77, p < 0.01) between activation during the first session and improvement in score. This negative correlation is a consequence of the facts that the correlation between mean activity and score on Session 2 (r=0.519) is stronger than on Session 10 (r=0.370) and that the Session 2 and Session 10 scores are strongly correlated (r=0.723).

The region we identified by correlating with overall score is not exactly the same as the Vo et al. region. Even if it were, there are enough other differences between their study and ours to make it hard to diagnose the source of our failure to find a positive relationship anywhere between striatal activation and improvement in score. We suspect that one of the reasons this study failed to find a significant relationship to learning was that early and late scores were so strongly correlated⁹. Our participants showed

⁹ Also the earlier studies used a joystick and its additional motor demands may

dramatic improvements over the course of the experiment, but the best performers tended to stay the best performers. Their performance was largely a function of the strategy they adopted early in playing the game and the effect of practice was largely to hone that strategy.

5. Conclusions

This research has indicated that one can predict individual differences in performance by classifying scans of brain activity into discrete states and then examining the sequential structure of these states. The Space Fortress game allows a direct approach to the classification step because we can map external states of the game directly onto model states. This then provides an objective basis for training a classifier to recognize the states. This is possible in Space Fortress because the internal state of the ACT-R model (and probably any successful cognitive model of the task) is highly reactive to the external state of the game. Such a direct mapping between external states and relevant mental states is true for many human activities (such a driving) but hardly all. In more reflective situations like mathematical problem solving people go through critical mental states that are not indicated by changes in the external world or in their overt behavior.

We have also modeled such reflective tasks with hidden Markov models (HMMs - because the states are now hidden) where we infer people's mental states by changes in brain activation and connect these HMM models to ACT-R cognitive models (e.g. Anderson and Fincham, 2014). These efforts differ from the current effort in that there is not an external definition of the mental state. Rather, we induce a mapping from brain activation to model states that maximizes the regularity in the imaging data. We went back to the data in Anderson & Fincham, classified the scans into states given the induced mapping, and then looked at the log probabilities of these scans according to the model described in that paper. We did find a positive correlation (r=0.644 for 75 participants) between this measure and percent of the problems participants got correct. Thus, it may be possible more generally to use the sequential structure of brain activation to predict success in a task.

Acknowledgments

This work was supported ONR grant N000140910098, and a James S. McDonnell Scholar Award. We would like to thank Aryn Pyke for her comments on the manuscript. Correspondence concerning this article should be addressed to John R Anderson, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213 (email: ja@cmu.edu). The analyses and models in this paper can be obtained at http://act-r.psy.cmu.edu/?post_type=publications&p=16893

Appendix A. Regions showing Effects of Binary Contrasts and Correlation with Score

We performed a standard GLM of bold activity to determine which regions showed effects of four binary variables – whether it was Session 2 or Session 10, whether the fortress was present or not, whether the mine was present or not, and whether or not the participant was in the top half of the players. The threshold for

(*footnote continued*) somehow be critical.

significance for the interaction was at least 10 contiguous voxels at a voxel-wise alpha of p < 0.0001. Spatial smoothness of these data were computed over the residuals and used as input for simulation to estimate brain-wise alpha to be less than 0.01 (Cox, 1996; Cox and Hyde, 1997).

No region showed a significant difference between better and weaker players. Table A1 lists regions showing significant effects of the other binary contrasts. Activation was higher in Session 2 relative to Session 10 in superior parietal regions associated with spatial attention and the fusiform areas, suggesting that participants are better managing the visual demands of the task at the end of the experiment. The fortress conditions show increased activation in the left motor region controlling the right hand, and in the bilateral fusiform region. This reflects the demand for shooting and monitoring of number of shots. The mine conditions show increased activity in the lateral inferior prefrontal cortex and the posterior parietal cortex, which are regions in the ACT-R model that are responsible for handling the memory demands of the mines conditions.

We also conducted an analysis of the correlation between average score over the two sessions and average voxel activation. Using the same threshold as with the binary contrasts (10 contiguous voxels with a significance less than 0.0001) no region reached significance. However, if we used an alternative threshold more appropriate for correlation of p=0.0025 with 40 contiguous voxels, which also corresponds to a brain-wise alpha of 0.01, we did find one right dorsal striatum region that consisted of 52 voxels overlapping with the caudate and nearby regions like the claustrum (Talairach coordinates of x = 21, y = 7, z = 22).

Training in all versus specific conditions

Fig. 7 presents classification performance training the classifier on all other games but the game being predicted. So for instance, in predicting the Navigation states for a Fortress-Only game we are also including games from different experimental conditions such as Mines-Only. This raises the question of whether it would be better to restrict training just to games from the same condition. Table A2 compares success in prediction when training and testing on different combinations of conditions.

The training cases consist of all games from the 35 other participants plus 35 copies of the games of the participant. In the cases where the training and test conditions are the same, we excluded the game being tested in making these 35 copies for training. The measure reported in the table is the D-prime measure of discriminability (Wickens, 2002) generalized to four categories. This involves calculating 4 D-primes, one for each state. The D-prime for a state is calculated from probability of a "hit", which is correctly labeling a scan from that state, and the probability of a "false alarm", which is assigning a scan from a different state to that state. The definition of D-prime is

d - prime = Z(p - hit) - Z(p - false - alarm)

where Z is the inverse of the cumulative Gaussian distribution. The D-primes reported in the table are the average of the D-primes for the 4 states. The D-primes are further averaged over participants and Sessions 2 and 10. Parts (a) – (c) of the table report performance for the three groups of states. All the D-primes in Table A2 are positive and they are positive for all participants. Thus, with all combinations of training and testing categories there is information to discriminate states. Part (d) of the table reports results averaged over different experimental conditions. Testing on the same condition as trained results in better performance than training on a different condition (second versus first row in part d).

Table A1

Regions showing Significant Effects of Binary Variables.

		Voxel count	Brodmann Areas	Coordinates (x, y, z)	Effect				
(a) Session 2 : Session 10									
1	R. Precuneus, Superior Parietal Lobule	181	7,19	8,-62,50 >					
2	L. Precuneus, Middle Occipital G, Cuneus	80	40	-25, -78,30 >					
3	L. Postcentral G, Inferior Parietal Lobule	17	9,6	-59,-20,26 >					
4	L. Middle/Inferior Temporal G, Middle Occipital G	143	46,10,9	-47,-63,2 >					
(b) Fc	ortress : No Fortress								
1	L. Precentral/Postcentral G, Inferior/Superior Parietal Lobules, Middle Frontal G	1226	6,4,3,40	-37,-22,53 >					
2	M. Precuneus	64	7	0,-51,49 >					
3	M. Medial Frontal G, Cingulate G	644	6,32,24	1,2,48 >					
4	R. Precuneus, Superior Parietal Lobule	59	7	27, -48, 48 >					
5	R. Middle Frontal G, Precentral G	407	6	36, -3, 47 >					
6	L. Inferior Parietal Lobule, Angular G, Supramarginal G	121	39, 40	-45, -59,37 <					
7	R. Middle/Superior Frontal G	368	9, 8	28,37,30 >					
8	L. Middle/Superior Frontal G	66	9	-29,37,30 >					
9	L. Inferior Frontal G, Precentral G	17	6,9	-37,4,28 >					
10	M. Cuneus, Precuneus	73	18,7,31	-3,-75,27 <					
11	L. Precuneus, Posterior Cingulate	19	31	-12,-51,26 <					
12	L. Postcentral G. Inferior Parietal Lobule	39	2	-55, -18, 24 >					
13	L. Insula	21	13	-29,21,10 >					
14	L/R. Thalamus	387		-7,-15,9 >					
15	R. Insula, Inferior Frontal G	64	13	32,23,8 >					
16	R. Middle/Inferior Temporal G, Middle/Inferior Occipital G	166	37,19	45,-67,5 >					
17	L. Middle Occipital G, Middle Temporal G, Declive	389	19,37,18	-39, -70,0 >					
18	L. Fusiform G, Culmen	29	37,19	-29, -53, -11 >					
19	R. Culmen, Declive, Fusiform G	271	37,19	18, -52, -12 >					
20	L. Parahippocampal G, Amygdala, Uncus, Superior Temporal G	70		-27, -3, -19 <					
(c) Mines									
1	R. Middle/Superior Frontal G	44	6	31, -4,61 >					
2	L. Precentral/Postcentral G, Middle Frontal G	363	6,4,3,40	-37, -16,53 >					
3	M. Medial Frontal G, Cingulate G	166	6,32,24	-2,6,45 >					
4	R. Middle Frontal G	43	8,9	32,34,41 >					
5	L/R. Precuneus, Cuneus	699	7,19	-13,-64,40 >					
6	L. Inferior/Middle Frontal G, Precentral G	113	9,6	-41,5,29 >					
7	R. Middle Temporal G	17	19	45,-63,15 >					
8	L. Insula, Middle Frontal G	112	13	-32,24,13 >					
9	L. Thalamus	106		– 12, – 12,13 >					
10	L/R. Lingual G, Cuneus, Posterior Cingulate, Declive	1119	18,19,30,17	-1,-72,2 >					
11	L. Fusiform G, Middle Occipital G	71	37,19	-38, -56, -7 >					

Table A2

D-prime measures of discriminability given different training sets

	Test Sets					
		Orbit	Fortress- Only	Mines- Only	Both	
(a)Training Sets for Navigation	Orbit Fortress- Only	0.834 0.679	0.637 0.917	0.521 0.614	0.334 0.660	
	Mines- Only	0.617	0.658	0.664	0.501	
	Both	0.543	0.848	0.619	0.686	
	All	0.784	0.933	0.719	0.718	
				Test Sets		
			Fortress-	Both		
			Only			
(b) Training Sets for	Fortress	Fortress-Only	0.920	0.767		
		Both	0.622	0.617		
		All	0.965	0.711		
			Test Sets			
			Mines-	Both		
			Only			
(c) Training Sets for Mines		Mines-Only	1.110	0.869		
		Both	0.633	0.900		
		All	1.094	0.928		
		Types of States				
	-	Navigation	Fortress	Mines		
(d) Averaged over	Different	0.603	0.695	0.751		
Conditions	Same	0.775	0.768	1.005		
	All	0.788	0.838	1.011		

In all cases these differences are highly significant (smallest effect is t(35) = 3.67, p < 0.001 for Fortress states). Training on all conditions (third line in part d and what was used for Fig. 7) results in slightly better performance than training on just the same condition. The difference is significant only for Fortress states, but there it is highly significant (t(35) = 5.49, p < 0.0001).

In summary, there is a remarkable generalization of features discriminating states in one experimental condition to another experimental condition. While there is discriminating information unique to a condition, it is slightly better to aggregate all conditions, taking advantage of the larger database and variety of training instances.

The deconvolution step

The deconvolution step is an attempt to infer the activity during a scan from the pattern of BOLD activity in surrounding scans (largely focusing on later scans). As noted in the paper we use a Weiner filter assuming the canonical SPM shape for the hemodynamic function. Handwerker et al., (2004) have argued that there is considerable variation in the actual function across regions and participants. They estimate ranges in time to peak from 2.5 to 6.5 whereas the SPM function has a time to peak of 5 s. We explored the performance of our model, using the standard double gamma but setting the parameters so that time to peak varied in 1-second increments from 3 to 7 s. Fig. A1a shows the correlation of the activity inferred for a scan with the BOLD signal on subsequent scans assuming hemodynamic responses with different peaks. As can be seen, assuming the SPM function with a 5 s peak the maximum correlation is with the third scan, which makes sense since this scan spans the 4–6 s. In some of our other work, we have just used the BOLD activity from two 2-second scans later to infer activity during a scan. Similarly, the other peaks tend to have maximum correlation with the BOLD response on the scans that contain the peak.

A question of interest is whether the SPM function is the best choice of a function and also whether we actually do better going through the deconvolution step rather than just using BOLD activity at a later point. To address this question Fig. A1b plots the average D-prime classification of scan labels using functions with different peaks and Fig. A1c plots the average classification using different scans to classify activity on scan 1. The best classification is obtained with the standard SPM peaks of 5 seconds but this is not superior to 6 s (t(35)= 1.37, p < 0.05). However, it is superior



Fig. A1. (a) Correlation of the output of a Weiner filter for a scan in our data with the BOLD response during subsequent scans. The different lines are the results when hemodynamic responses are used that have different peaks. (b) D-prime measures of accuracy of state classification using different peaked hemodynamic responses in a Weiner filter. (c) D-prime measures of accuracy of state classification using the BOLD response during different scans.

to 3 or 4 s for all 36 subjects and superior to 7 s for all but 2 subjects. As can be seen from Fig. A1c using the bold activity 2 scans later is clearly superior (for all subjects in the case of scans 0, 1, and 4 scans later and for 32 of 36 in the case of 3 scans later). While using the BOLD signal 2 scans later results in a D-prime of 0.89, using deconvolution with the 5 s SPM response is better resulting in a D-prime of 0.94. This small difference holds for 30 of he 36 subjects and results in a highly significant t=5.28, p < 0.0001.

Explanation of semi-Markov structure in Fig. 5

State Transitions (Part a): Initially, there is no mine and so the model starts in the None state. If a friendly¹⁰ mine appears the player first shoots it and then corrects the orbit of the ship, which has drifted from the circular path while the mine is being processed. If a Foe mine appears it must first be IFF'd (neutralized) before it can be shot. This involves making a double key-press between 250 and 400 m s. Randomly half of the mines are Friendly and half are Foe, resulting in an equal probability of transitioning from a None scan to either a Shoot scan or an IFF scan. Even though the orbit usually needs to be fixed after shooting the mine, there is only 35% probability of transitioning from a Shoot scan to a Fix scan because 65% of the time the orbit can be fixed in the same scan as the shooting and so the next scan is in the None state. Similarly, even though a Foe mine still needs to be shot after being IFF'd, there is only 50% probability that the next scan is a Shoot scan. 29% of the time the mine is shot in the same scan as the IFF and the next scan is in the Fix state. 21% of the time the mine is shot and the orbit fixed in the IFF scan, and the next scan is in the None state. Usually (74%), after the orbit is fixed the next scan is in the None state because no mine has yet appeared. However, sometimes a mine will appear in the next scan and that next scan will then be either in the Shoot or IFF state.

State Durations (part b): More successive scans occur in the None state than any other state. The actual duration of these "rest" periods is determined both by the random appearance of mines and how fast the player can dispose of the mines. Frequently, the period of time from which an IFF can be executed to when it is completed overlaps two scans but sometimes it can be completed within a single scan. Shooting a mine requires aiming and waiting until the mine appears at the aim point. Like an IFF, this is usually either completed in one scan or overlaps two scans. The corrections to orbits are usually brief and can be completed in a single scan.

Navigation and Fortress Spaces

Fig. 8 provided a two-dimensional projection of the Mine states onto a two dimensional plane that captured most of the variance among the states. Fig. A2 provides the same information for the Navigation and Fortress states. The two-dimensional structure for Navigation in part (a) accounts for 84% of the variance and the two-dimensional structure for Fortress in part (b) accounts for 71% of the variance.

¹⁰ The terms "Friendly" and "Foe", which are part of game, can be misleading. A Friendly mine can still damage the ship and should be destroyed. The only difference is that one does not have to execute the double-click to neutralize a Friendly mine. Indeed, doing so will energize a Friendly mine and make it invulnerable.



Fig. A2. Projection of the 4 Navigation states (a) and the 4 Fortress States (b) onto 2 dimensions. The origin of this space corresponds to the average activation during Mine games. The X and Y vectors indicate the activity added to this average activation for individuals types of scans.

References

- Anderson, J.R., Farrell, R., Sauers, R., 1984. Learning to Program in LISP1. Cognit. Sci. 8 (2), 87–129.
- Anderson, J.R., 2005. Human symbol manipulation within an integrated cognitive architecture. Cognit. Sci. 29 (3), 313–341.
- Anderson, J.R., 2007. How Can the Human Mind Occur in the Physical Universe?. Oxford University Press, New York.
- Anderson, J.R., Betts, S., Ferris, J.L., Fincham, J.M., 2010. Neural Imaging to Track Mental States while Using an Intelligent Tutoring System. Proc. Natl. Acad. Sci. USA 107, 7018–7023.
- Anderson, J.R., Fincham, J.M., Yang, J., Schneider, D.W., 2012. Using brain imaging to track problem solving in a complex state space. NeuroImage 60, 633–643.
- Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., Qin, Y., 2004. An integrated theory of mind. Psychol. Rev. 111, 1036–1060.
- Anderson, J.R., Bothell, D., Fincham, J.M., Anderson, A.R., Poole, B., Qin, Y., 2011. Brain regions engaged by part-and whole-task performance in a video game: a model-based test of the decomposition hypothesis. J. Cognit. Neurosci. 23 (12), 3983–3997.
- Anderson, J.R., Fincham, J.M., 2014. Extending problem-solving procedures through reflection. Cognit. Psychol. 74, 1–34.
- Bilalić, M., Langner, R., Ulrich, R., Grodd, W., 2011. Many faces of expertise: fusiform face area in chess experts and novices. J. Neurosci. 31 (28), 10206–10214.
- Birn, R.M., Smith, M.A., Jones, T.B., Bandettini, P.A., 2008. The Respiration response function: the temporal dynamics of fMRI signal fluctuations related to changes in respiration. NeuroImage 40, 644–654.
- Bothell, D., 2010. Modeling Space Fotress: CMU Effort. (http://act-r.psy.cmu.edu/ workshops/workshop-2010/schedule.html).
- Chein, J.M., Schneider, W., 2005. Neuroimaging studies of practice-related change: fMRI and meta-analytic evidence of a domain-general control network for learning. Cognit. Brain Res. 25 (3), 607–623.
- Chein, J.M., Schneider, W., 2012. The brain's learning and control architecture. Curr. Dir. Psychol. Sci. 21 (2), 78–84.
- Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput. Biomed. Res., Int. J. 29, 162–173.
- Cox, R.W., Hyde, J.S., 1997. Software tools for analysis and visualization of fMRI data. NMR Biomed. **10**, 171–178.
- Destefano, M., Gray, W.D., 2008. Progress report on Pygame Space Fortress. Rensselaer Polytechnic Institute, Troy, NY.
- Donchin, E., 1989. The learning-strategies project: introductory remarks. Acta Psychol. 71 (1-3), 1–15.
- Erickson, K.I., Boot, W.R., Basak, C., Neider, M.B., Prakash, R.S., Voss, M.W., Kramer, A.F., 2010. Striatal volume predicts level of video game skill acquisition. Cereb. Cortex 20 (11), 2522–2530.

- Frederiksen, J.R., White, B.Y., 1989. An approach to training based upon principled task decomposition. Acta Psychol. 71 (1–3), 89–146.
- Friston, K.J., Ashburner, J.T., Kiebel, S.J., Nichols, T.E., Penny, W.D. (Eds.), 2011. Statistical Parametric Mapping: The Analysis of Functional Brain Images: The Analysis of Functional Brain Images. Academic Press, United States.
- Glover, G.H., 1999. Deconvolution of impulse response in event-related BOLD fMRI. NeuroImage 9, 416–429.
- Gopher, D., Weil, M., Siegel, D., 1989. Practice under changing priorities: an approach to training of complex skills. Acta Psychol. 71, 147–178.
- Gobel, E.W., Parrish, T.B., Reber, P.J., 2011. Neural correlates of skill acquisition: decreased cortical activity during a serial interception sequence learning task. Neuroimage 58 (4), 1150–1157.
- Handwerker, D., Ollinger, J., D'Esposito, M., 2004. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. NeuroImage 21 (4), 1639–1651.
- Haxby, J.V., Hoffman, E.A., Gobbini, M.I., 2000. The distributed human neural system for face perception. Trends Cognit. Sci. 4 (6), 223–233.
- Larkin, J.H., 1981. Enriching formal knowledge: a model for learning to solve textbook physics problems. In: Anderson, J.R. (Ed.), Cognitive skills and their acquisition. pp. 311–334.
- Nestor, A., Plaut, D.C., Behrmann, M., 2011. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. Proc. Natl. Acad. Sci. 108 (24), 9998–10003.
- Norman, K.A., Polyn, S.M., Detre, G.J., Haxby, J.V., 2006. Beyond mind-reading:
- Multi-voxel pattern analysis of fMRI data. Trends Cognit. Sci. 10, 424–430. Pereira, F., Mitchell, T., Botvinick, M., 2009. Machine learning classifiers and fMRI: a
- tutorial overview. Neuroimage 45, S199–S209. Poldrack, R.A., Desmond, J.E., Glover, G.H., Gabrieli, J.D., 1998. The neural basis of
- visual skill learning: an fMRI study of mirror reading. Cereb. Cortex 8 (1), 1–10. Rabiner, L.R., 1989. A tutorial on hidden Markov models and selected applications
- in speech recognition. Proc. IEEE 77 (2), 257–286. Scherf, K.S., Luna, B., Avidan, G., Behrmann, M., 2011. "What" precedes "which":
- Developmental neural tuning in face-and place-related cortex. Cereb. Cortex 21 (9), 1963–1980.
- Steiger, J.H., 1980. Tests for comparing elements of a correlation matrix. Psychol. Bull. 87 (2), 245.
- Tarr, M.J., Gauthier, I., 2000. FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. Nat. Neurosci. 3, 764–770.
- Vo, L.T., Walther, D.B., Kramer, A.F., Erickson, K.I., Boot, W.R., Voss, M.W., Wang, M. Y., 2011. Predicting individuals' learning success from patterns of pre-learning MRI activity. Plos One 6 (1), e16093.
- Wickens, T.D., 2002. Elementary Signal Detection Theory. Oxford University Press, New York.
- Yu, S.Z., 2010. Hidden semi-Markov models. Artif. Intell. 174 (2), 215-243.