Individual Differences in Decision Making Strategies Can be Predicted by Resting-State Functional Connectivity

Yuxue Cher Yang (chery@uw.edu)
Department of Psychology, University of Washington, Seattle, WA 98195 USA

Catherine Sibert (sibert@uw.edu)
Department of Psychology, University of Washington, Seattle, WA 98195 USA

Andrea Stocco (stocco@uw.edu)
Department of Psychology, University of Washington, Seattle, WA 98195 USA

Abstract

As the study of individual differences becomes more widespread, questions arise about the reasons that a particular individual might adopt a particular strategy. Using both the behavioral and functional neuroimaging data of healthy adults from Human Connectome Project (HCP) we examined decision making in an incentive processing task (Delgado et al. 2000). A pair of distinct ACT-R models, representing a Declarative strategy and a Procedural strategy, were used to classify subjects as either Declarative or Procedural decision makers based on their behavioral data. A machine learning Lasso analysis was performed on each subject's resting state functional connectivity, and was able to match the ACT-R model classifications to a high degree of accuracy. This suggests that the strength of connections between brain regions may play an important role in shaping the decision making process of a given individual.

Keywords: Decision Making; Strategy; Computational modeling; Functional connectivity; Procedural Memory; Declarative Memory; ACT-R

Introduction

It has been argued that, to be effective, computational cognitive models need to switch from nomothetic, group-level descriptions to idiographic, individual-level ones (Zhou et al., 2021). A promising framework in this sense was proposed by Ritter and Gobet (2000), who argued that an architecture can be used to successfully capture the invariant part of the mind, while different parameter values can be used to model variations across individuals. This approach was tested successfully by Daily and Lovett (2001), who succeeded in capturing individual differences in working memory through a single parameter in the ACT-R architecture (spreading activation *W*), and more recently, by Xu and Stocco (2021) using behavioral data. Recent work has also shown that individual parameter values are associated with different signatures of neural activity in EEG data (Zhou et al., 2021) and fMRI (Rice & Stocco, in press). These neural signatures were identified from "resting-state" recordings, that is, task-free sessions in which participants are not asked to do anything in particular, and which offer the opportunity to observe spontaneous but highly organized brain activity (Fox et al., 2005). The fact that parameter values that capture individual differences are reflected in resting state imaging data suggests a biological underpinning for these parameters.

Despite its successes, the approach of identifying individual differences with parameter values still runs into conceptual roadblocks. While a cognitive architecture can be assumed to reflect an invariant, innate blueprint (Taatgen, 2020), participants are typically measured when performing a *specific task*, and, even with the same architecture, participants might perform the same tasks in the same way. For example, simple association learning tasks can be modeled using two strategies, a procedural-based reinforcement learning strategy and a memory-based, decision-by-sampling or instance-based learning strategy. Haile et al. (2020) showed that different participants are best fit by different strategies. This implies that attempts to measure single parameters across participants is ultimately doomed to fail: it does not make sense to estimate learning rate (a reinforcement learning parameter) from participants who rely on memory, and it does not make sense to measure rate of forgetting (a successfully decodable parameter) from individuals who follow a memory-less, procedural learning strategy.

Through computational models, it is possible to make inferences about which strategy a participant is using (Haile et al., 2020). But what makes participants *prefer* a strategy over another? In principle, strategy selection could be a function of personal preference, habit, or cost-benefit analysis (Payne, Bettman, & Johnson, 1993). One enticing possibility is that strategy selection might reflect bounded rationality (Lewis et al., 2014): individuals choose the strategy that plays to their strengths, yielding the best results given the computational costs involved. If this is the case, then it follows that preference for a strategy over another would also ultimately depend on identifiable stable characteristics of their brain activity.

To test this hypothesis, we analyzed a dataset including almost 200 participants for whom performance on a simple decision-making task and resting-state fMRI data were available. Computational models implementing alternative strategies were fit to individual behavioral data to determine the most likely strategy used by each participant. Machine learning techniques were then employed to identify the facets of spontaneous neural activity that best predict which strategy will be used by each individual. We expected to find that decision-making strategies associated with the use of memory resources (such as retrieving the previous success history of an option) would be associated with increased functional connectivity in fronto-parietal regions responsible for cognitive control. Conversely, we expected that decision-making strategies associated with habitual and reward-based learning would be associated with increased functional connectivity in sensorimotor cortices responsible for automatic stimulus-response behaviors and with the basal ganglia circuit responsible for feedback-driven learning (Yin & Knowlton, 2006).

Methods

This study analyzed both behavioral and neuroimaging data obtained from the Human Connectome Project (HCP) dataset (Van Essen et al., 2013). A total of 199 participants (111 females, 85 males, and 3 did not disclose) who completed both sessions of the task-based fMRI gambling game were included in this study. All participants were healthy adults with no neurodevelopmental or neuropsychiatric disorders. The experimental protocol, subject recruitment procedures, and consent to share de-identified information were approved by the Institutional Review Board at Washington University.

The Incentive Processing Task in the HCP

This incentive decision making task was adapted from the gambling paradigm developed by Delgado and Fiez (2000). Participants were asked to guess if the number on a mystery card (represented by a "?", and ranging from 1-9) was more or less than 5. After making a guess, participants were given feedback, which could take one of three forms, *Reward* (a green up arrow and $1), *Loss* (a red down arrow and -$0.50), or *Neutral* (a gray double headed arrow and the number 5). The feedback did not depend on the subject's response, but was determined in advance; the sequence of pre-defined feedback was identical for all participants. The task was presented in two runs, each of which contains 64 trials divided into eight blocks. Blocks could be *Mostly Loss* (6 loss trials pseudo-randomly interleaved with either 1 neutral and 1 reward trial, 2 neutral trials, or 2 reward trials) or *Mostly Reward* (6 reward trials pseudo randomly interleaved with either 1 neutral and 1 loss trial, 2 neutral trials, or 2 loss trials). In each of the two runs, there were two Mostly Reward and two Mostly Loss blocks, interleaved with 4 fixation blocks (15 seconds each).

Resting-State fMRI Analysis

This study employed the "minimally preprocessed" version of resting-state fMRI data, which has already undergone a minimal number of standard preprocessing steps including artifact removal, motion correction, normalization, and registration to the standard MNI ICBM152 template. Additional preprocessing steps were performed using the AFNI software (Cox RW, 1996), including despiking, spatial smoothing with an isotropic Gaussian 3D filter FWHM of 8 mm, and removal of linear components related to the six motion parameters and their first-order derivatives.

Functional connectivity measures were constructed from the HCP resting-state data using Power et al. (2011)'s whole brain parcellation. This parcellation was used to construct a 264 Region of Interest (ROI) functional atlas, with each ROI containing 81 voxels. This parcellation atlas is defined in the MNI space and was applied to all participants in HCP dataset. The extraction of the time series and calculation of the connectivity matrices was performed using R (RStudio Team, 2016) and Python. Pearson correlation coefficients and partial correlation coefficients between the time series of each brain region were calculated for each participant, resulting in a 264×264 symmetric connectivity matrix for each session for each subject. The averaged correlation coefficients across subjects were calculated by first transforming each r value into a z-value, and then retransforming the average z value back into an equivalent r value using the hyperbolic tangent transformation (Silver & Dunlap, 1987).

Response Switch Analysis

Because in the Incentive Processing task the feedback is scheduled in advance and does not depend on actions taken by participants, it is impossible to define participant's performance in terms of either accuracy or learning. This poses a challenge when trying to determine if participants are responding to feedback. The most meaningful way to check whether participants change their behavior in response to feedback is through analyzing their Win-Stay, Lose-Shift (WSLS) probabilities. Thus, our main dependent variable was the tendency to switch responses after a Loss feedback and after a Reward feedback. This response switch is coded as 0 if the current response is the same as the next response, and coded as 1 if the current response is not the same as the next response. Because the response switch is a binary variable, the analysis was conducted with logistic mixed-effects models using orthogonal contrast coding as implemented in the "lme4" package in R. Given that Neutral trials make up only a small proportion of total trials, they were excluded in statistical tests. In the mixed-effect model, Block Type (Reward or Loss) and Trial Type (Reward or Loss) were treated as fixed effects, and individual participants were treated as random effects. The parameters were estimated based on the maximum likelihood.

On the group-level, there is no significant effect of feedback nor Block Type on the probability of switching responses. However, and critically for this study, on an individual-level, individuals do exhibit different behavioral response profiles. Figure 1 demonstrates the mean probability of response switching as a function of Trial Type (feedback received) and Block Type.

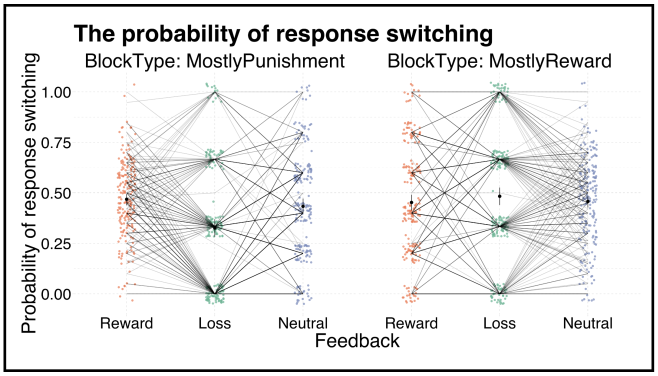

Figure 1: The mean probability of response switching as a function of feedback and block type. Each color dot and grey line represents the mean probability of response switching of a single participant, and the black dot represents the mean and 95% confidence interval across participants.

We also examined whether the response times change as a function of previous feedback (the Trial Type of the previous trial) and Block Type. Excluding neutral trials in the statistical analysis, on average, participants tend to take longer when making decisions in Mostly Reward blocks than in Mostly Loss blocks (β= 15.21, *SE* = 6.39, *p* = 0.017), regardless of previous feedback. Figure 2 shows the mean response time (RT) as a function of Previous Feedback and Block Type. Compared to the probability of response switching, however, the pattern of RTs was found to be noisier and less consistent across individuals, and was therefore not included in the following modeling analysis.

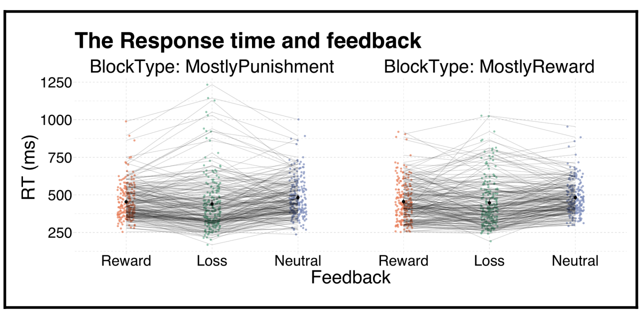

Figure 2: Mean response time as a function of previous trial feedback and block type. Each colored dot and grey line represents the RT of an individual, and the black dot represents the mean and confidence interval (95%) of RT across participants.

ACT-R Model Design

While the behavioral data does not reveal major effects across subjects, it offers an exciting opportunity from a modeling perspective. There exist two competing explanations of how decision making occurs in a repeated choice paradigm, one based on episodic memory of previous choices (Gonzalez et al., 2003) and one based on reinforcement learning (Daw et al., 2011). Each explanation is dependent on different mechanisms, and, ultimately, reliant on different strategies. Both explanations were implemented as two computational models in the ACT-R cognitive architecture (Anderson, 2007): as a *Declarative* Model, reliant on memory retrieval, and a *Procedural* Model, which makes use of reinforcement learning.

Declarative Model The Declarative Model relies on the declarative module to retrieve a memory of prior actions and their corresponding feedback. When presented with a mystery card, the model selects an action, LESS or MORE, for evaluation, and makes a retrieval of the prior history of feedback associated with that action. If the retrieved history contains a WIN result of the chosen action, it will execute that action, but if the history contains a LOSE or NEUTRAL result, the model will execute the alternate action. If no history is retrieved, an action will be executed at random. After making a guess, the model is presented with feedback, which is encoded as a new memory chunk associated with the selected action. In ACT-R, memory chunks are retrieved based on their activation, calculated with a base-level learning function that reflects the degree to which a chunk matches the context of the retrieval request, and the recency of prior retrievals (Eq 1). If the activation surpasses a specified threshold, the chunk is possible to retrieve and if multiple chunks meet this threshold, the chunk with the greatest activation will be selected. The model functions by remembering the results of previous actions to guide future actions.

$$A_i = \ln\left(\sum_{j=1}^{n} t_j^{-d}\right) + \epsilon \qquad (1)$$

Procedural Model By contrast, the Procedural Model represents the possible actions of the decision-making processes as competing rules, and reinforcement learning is used to increase the use of the rule that leads to the best outcomes. Instead of encoding each trial as a memory of an action and associated feedback, the model has two competing production rules that execute the MORE and LESS actions.When presented with the mystery card, the model will choose one of the rules to execute based on their utility. Initially, both rules have equal utility, and one will be chosen at random. After making a guess, the model is presented with a WIN, LOSE, or NEUTRAL response, and this feedback is encoded as an adjustment to the utility of the selected production rule (+1 for a WIN result, -1 for a LOSE result, and no change for a NEUTRAL result). At any time point *t*, the utility *U* of production *p* is calculated using Eq 2, where α indicates the learning rate, R_t is the reward the production received for at time *t*. Previous rewards will encourage the model to repeat the associated action, while a pattern of losses will decrease the utility of the action and encourage the selection of the alternate action.

$$U_t = U_{t-1} + \alpha(R_t - U_{t-1}) + s \qquad (2)$$

Individual Fit and Model Evaluation

To examine the predictions of our model, we used a grid-search approach to find the best possible parameters within the parameter space shown in Table 1. Each model simulates 64 trials, the same as the experimental paradigm for participants, repeated over 50 runs. The simulated stimuli were presented in the same order as the real experimental stimuli to avoid any potential noise from sequence effects in the simulation. Following the six conditions (Reward, Loss, Neutral trials in Mostly Reward Block and Reward, Loss, Neutral trials in Mostly Loss Block), the mean probability of response switching, P(Switch), and its standard deviation are computed.

Table 1: Model parameter space in the simulations.

Models	Parameter	Value	Meaning
Declarative	ϵ	0 - 0.5	activation noise
	d	0.2 - 0.85	memory decay
Procedural	s	0 - 0.5	utility noise
	α	0.05 - 0.5	learning rate

In order to evaluate the goodness-of-fit for individual fitting, we estimated maximum Log-Likelihood across the parameter space. The likelihood function of a particular model with parameters θ, $L(m,\ \theta|\ x)$, is the probability that, given the parameterized model and set of observed data to fit, the model would produce that data:= $L(m,\ \theta|\ x)$ = $P(x|m,\ \theta)$. Here, $m$and θ refers to the model and its parameters, and x refers to the observations. Common comparison metrics, such as the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC), are both based on likelihood, but rely on closed-form likelihood functions. While it is possible to derive such functions for simple models (such as logistic models or linear models), they can be incredibly difficult to derive for more complex models and impossible for arbitrarily complex models based on ACT-R and other high level architectures. Some attempts have been made to evaluate complex models with basic likelihood metrics: Stocco and Haile (2018), Prat and Stocco (2020), and Yang and Stocco (2019) have all used BIC to compare competing ACT-R models. However, the equation used to estimate BIC is a closed-form approximation that is based on Residual Sum of Squares and was originally derived for linear models; as such, it does not necessarily hold for ACT-R.

In this paper, we followed the computationally expensive but more accurate solution of empirically calculating the likelihood function by simulating each model and set of parameters multiple times, and calculating the empirical probability distribution of each set of results (Yang, Karmol, Stocco, in press). Knowing the mean and standard deviation of this distribution, the value of $P(x|\ m,\ \theta)$ can then be calculated directly. If a model is designed to predict n data points (corresponding, for instance, to different experimental conditions), its likelihood can be expressed as the joint probability that any of those data points can be produced. For simplicity, and assuming independence, this can be expressed as the product of the probability of observing each individual data point in the empirical data, i.e., $L(m,\ \theta\ |\ x_1,\ x_2,\ \dots\ x_n) = \prod_i L(m,\ \theta\ |\ x_i\ )$. Finally, to avoid computational problems with vanishing small probabilities, it is common the express this value in terms of *log* likelihood:

$$\log L = \log P(x\ |\ m,\ \theta) = \sum_i \log z[(x_i\ -\ x_{i,m})\ /\sigma_{i,m}]\ (3)$$

Results

Decision-Making Strategy Identification

By excluding participants who did not complete the gambling task and two sessions of resting state fMRI scanning, a total of 199 participants were fit by ACT-R models. Of these, 127 (63.82%) were best fit by the Declarative Model, and thus were identified as Declarative decision makers. The remaining 72 (36.18%) individuals were best fit by the Procedural Model, and identified as Procedural decision makers. The logistic mixed-effects model was conducted using orthogonal contrast coding as implemented in the lme4 package in R. ACT-R Model Type (Declarative vs. Procedural), Block Type (Mostly Reward vs. Mostly Loss), and Feedback (Reward vs. Loss) were treated as fixed effects, and individual subjects were treated as random effects. Full statistical results are shown in Table 2. In contrast to the lack of significant effects across the behavioral data, the probability of response switching was found to be statistically different between the two groups identified as either Declarative decision makers or Procedural decision makers (z = -6.11, p < 0.001), supporting the validity of the ACT-R model identification.

Table 2: Results of the Logistic Mixed Effects Model of the Probability of Response Switch

Statistical Test	**odds ratio**	*se*	*z*	*p*
(Intercept)	0.88*	0.05	-2.30	**0.022**
Model Group	0.71***	0.04	-6.11	**<0.001**
Block Type	0.98	0.03	-0.80	0.423
Trial Type	1.08**	0.03	2.84	**0.005**
Model Group by Block Type	1.07*	0.03	2.46	**0.014**
Model Group by Trial Type	0.8***	0.02	-8.10	**<0.001**
Block Type by Trial Type	1.02	0.03	0.78	0.434
Model Group by Block Type by Trial Type	1.10***	0.03	3.47	**0.001**
Random Effect				
σ^2	3.29			
ICC	0.12			

N HCPID	199
observation	9746
Marginal R^2 */Conditional* R^2	0.030/0.144
Log-Likelihood	-4010.323

* $p<0.05$ ** $p<0.01$ *** $p<0.001$

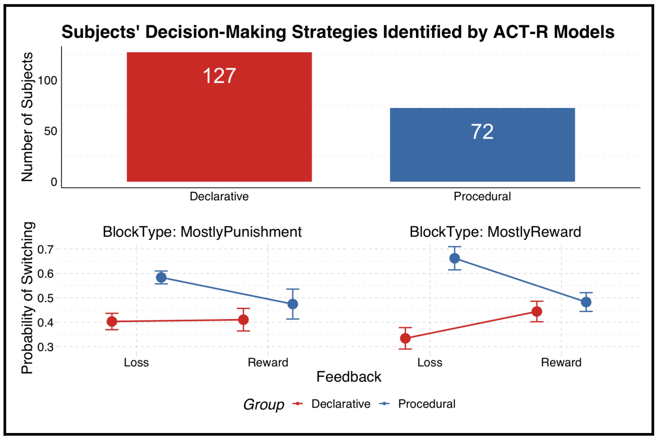

Figure 3. (Top) Counts of individuals identified by ACT-R models. The red bar represents the number of participants best fit by the Declarative model; the blue bar represents the number of participants best fit by the Procedural Model. (Bottom) The probability of response switching by two groups of individuals identified as either Declarative or Procedural decision makers.

Supervised Classification with Logistic Model

To explore if the behavioral differences between Declarative and Procedural decision makers are indicated by an individual's underlying brain structure, we trained a Logistic Regression model using resting state functional connectivity as its variable, and predicted the probability of a participant being labeled as either Declarative-based or Procedural-based decision maker by the ACT-R model classification. In order to handle an imbalanced dataset with unequal target labels, upsampling was applied by randomly adding data from the minority class. Having 69,696 (264 ROI × 264 ROI) connections, we want to select only the most important connections contributing to the prediction, therefore, Lasso regularization was applied to the Logistic Model. Lasso is a machine learning regression analysis technique that performs both variable selection and regularization in order to improve the prediction accuracy and interpretability of the computational model. It can reduce model complexity by penalizing large numbers of coefficients and also prevents overfitting which may result from simple linear regression. Lasso minimization is calculated using Eq 4, where the tuning parameter λ controls the degree of penalty: for greater values of λ, more coefficients are forced to become 0.

$$\sum_{i=1}^{n}\left(y_i - \sum_{j} x_{ij}\beta_j\right)^2 + \lambda \sum_{j=1}^{p} |\beta_j| \qquad (4)$$

To account for the large disparity between the number of participants and the number of predictors, we performed a Grid search in sklearn (Pedregosa et al., 2011) with 20-fold cross-validation to determine the best fit hyper-parameter λ (6.73). To maximize the penalty to coefficients, the highest value of λ with validation accuracy with one standard error of the maximum accuracy was chosen (as recommended by Krstajic et al (2014). Instead of splitting the entire dataset into training and testing sets, we refit the model using Leave-one-out (LOO) cross validation. The model is trained on all samples except one and the prediction is made on that one sample, then the process is repeated across the full dataset. The mean score (accuracy), true positive rate(TPR), true negative rate (TNR), false positive rate (FPR), false negative rate (FNR) are calculated across all folds to evaluate the performance of the model. By definition, the receiver operating characteristic curve (ROC) demonstrates the performance of a classification model by plotting the relationship between TPR vs. FPR at different classification thresholds. We calculated the AUC (Area under the curve), which is one of the most important metrics for evaluating a classification model's performance; as the AUC of a model approaches 1, the model approximates an ideal, perfect classifier. It provides information about how well a classification model is capable of distinguishing between classes. The overall classification accuracy is 0.88 and the ROC-AUC is 0.94, indicating that predicting from an individual's resting state functional connectivity, the Lasso Logistic model is capable of matching ACT-R's prediction about whether an individual is a Declarative-based or Procedural-based decision maker.

Connectivity Map

With Lasso regularization, approximately 1.4% of β estimates in the Logistic model are not zero, suggesting a relatively sparse neuro functional connectivity of the resting brain. In a 264 × 264 β coefficients matrix, the β_{ij} value indicates the weight of connectivity between the *i*-th and the *j*-th region in classifying whether the human subject is a Declarative decision maker or a Procedural decision maker from the resting state functional connectivity. The ultimate effect of β on the predicted group assignment depends on the polarity of the underlying functional connectivity. A positive β value has different implications if applied to a positive or negative partial correlation between two regions. To make the interpretation of the values unambiguous, we multiplied the βmatrix with the averaged partial correlation coefficient matrix A, obtaining a group-level weighted averaged correlation matrix W. Figure 4 demonstrates the brain connectivity map of *W*, thresholded so that only the most predictive 68 connections (corresponding to 0.01% of the initial pool of regressors) are shown. In this figure, red lines represent functional connections that are predictive of a Declarative decision maker, and blue lines represent functional connections that are predictive of a Procedural decision maker. Color shades suggest the strength of predictability.

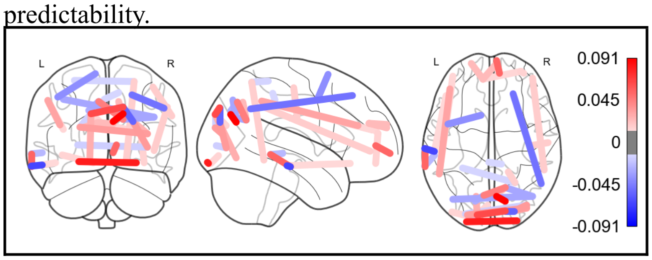

Figure 4. The group-level weighted averaged brain connectivity plot.

As we anticipated, the results show a dissociation between the types of connectivity associated with Declarative or Procedural strategies. Using the Power et al (2011) functional classification of these regions as a guideline, the results show that the use of a Declarative strategy was mostly associated with increased functional connectivity in the networks of regions associated with task control (fronto-parietal networks and attention networks) and episodic memory (default mode network and memory retrieval network), while the use of Procedural strategy was mostly linked to increased functional connectivity in sensorimotor and subcortical networks.

Discussion

This paper shows that individual preferences for using a declarative or a procedural strategy can be decoded from patterns of resting state functional connectivity data. The specific connectivity values suggest that an individual's preference for a particular strategy might be adaptive and rational. Specifically, individuals exhibiting a stronger fronto-parietal connectivity play to their strengths, and tend to use declarative strategies that are more reliant on controlled memory retrieval, while individuals with stronger sensorimotor connectivity tend to use procedural strategies. In general, the patterns of functional connectivity are compatible with ACT-R's regions.

Although our results are encouraging, a number of limitations must be acknowledged. First, the Declarative vs. Procedural classification of individuals' probability of switching is based on a log-likelihood model fitting procedure, and thus, no ground-truth labels were available. Moreover, the optimal parameter was searched from a finite grid, and determined by the highest log-likelihood value compared to empirical data. Second, the task is highly unusual, in that it provides no real opportunity for learning from feedback. Further study could model the learning effect and investigate whether different learning mechanisms could also be predicted by the neuro-functional connectivity.

These limitations notwithstanding, we believe that our results have some important implications. First, they provide a new and deeper way to connect individual differences in task performance with individual neurobiology, showing how the latter might provide constraints on the specific strategies that are selected.

Second, they have implications for ACT-R. Procedural knowledge has been traditionally associated, in ACT-R, with the function of the basal ganglia. While the role of the basal ganglia in learning procedural knowledge is well supported (Knowlton et al., 2006 etc.), it is not clear that the basal ganglia are also the ultimate seat of procedural knowledge. In fact, both modeling work (Stocco, Lebiere, & Anderson, 2010) and experimental work using neurostimulation (Rice & Stocco, 2019) point to procedural knowledge being ultimately encoded in a set of cortico-cortical connections that directly link stimulus-response associations. This interpretation is compatible with our findings that find greater likelihood of using procedural strategies in individuals with stronger perceptuo-motor connectivity.

Acknowledgements

This research was supported by grant FA9550-19-1-0299 from the Air Force Office of Scientific Research to AS.

References

Anderson, D. R. (2007). Model based inference in the life sciences: a primer on evidence. *Springer Science & Business Media*.

Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research, 29*(3), 162-173.

Daily, L. Z., Lovett, M. C., & Reder, L. M. (2001). Modeling individual differences in working memory performance: A source activation account. *Cognitive Science, 25*(3), 315-353.

Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience, 8*(4), 429-453.

Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., & Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of neurophysiology, 84*(6), 3072-3077.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences, 102*(27), 9673-9678.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance‐based learning in dynamic decision making'. *Cognitive Science, 27*(4), 591-635.

Haile, T., Prat, C. S., & Stocco, A. (2020) One size doesn't fit all: Idiographic computational models reveal individual differences in learning and meta-learning strategies. *Proceedings of the 18th International Conference on Cognitive Modeling*, pp. 75-81.

Lewis, R. L., Howes, A., & Singh, S. (2014). Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science*, 6(2), 279-311.

Krstajic, D., Buturovic, L. J., Leahy, D. E., & Thomas, S. (2014). Cross-validation pitfalls when selecting and assessing regression and classification models. *Journal of Cheminformatics*, 6(1), 1-15.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The Adaptive Decision Maker*. Cambridge university press.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *The Journal of machine Learning research, 12,* 2825-2830.

Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., ... & Petersen, S. E. (2011). Functional network organization of the human brain. *Neuron, 72*(4), 665-678.

Rice, P., & Stocco, A. (2019). The role of dorsal premotor cortex in resolving abstract motor rules: Converging evidence from Transcranial Magnetic Stimulation and cognitive modeling. *Topics in cognitive science, 11*(1), 240-260.

Rice, P., & Stocco, A. (in press). Estimating individual differences in working memory through ACT-R modeling and resting state connectivity. *Proceedings of the 19th International Conference on Cognitive Modeling.*

Silver, N. C., & Dunlap, W. P. (1987). Averaging correlation coefficients: should Fisher's z transformation be used?. *Journal of applied psychology, 72*(1), 146.

Stocco, A., Lebiere, C., & Anderson, J. R. (2010). Conditional routing of information to the cortex: A model of the basal ganglia's role in cognitive coordination. *Psychological review*, *117*(2), 541.

Taatgen, N. A. (2020). Cognitive Architectures: Innate or Learned?. *Common Model of Cognition Bulletin, 1*(2), 476-480.

The R Team (2016). *RStudio: integrated development environment for R*. RStudio, Inc., Boston, Massachusetts, USA.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., Ugurbil, K., & Wu-Minn HCP Consortium. (2013). The WU-Minn human connectome project: an overview. *Neuroimage, 80*, 62-79.

Xu, Y., & Stocco, A. (2021). Recovering reliable idiographic biological parameters from noisy behavioral data: The case of basal ganglia indices in the Probabilistic Selection Task. *Computational Brain & Behavior*, 1-17.

Yang, Y.C. and Stocco, A. (2019). Syntactic priming depends on procedural, reward-based computations: Evidence from experimental data and a computational model. *Proceedings of the 17th International Conference on Cognitive Modeling,* pp. 307-313.

Yang, Y. C., Karmol, A. M., & Stocco, A. (in press). Core cognitive mechanisms underlying syntactic priming: A comparison of three alternative models. *Frontiers in Psychology.*

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience, 7*(6), 464-476.

Zhou, P., Sense, F., Van Rijn, H., & Stocco, A. (2021). Reflections of idiographic long-term memory characteristics in resting-state neuroimaging data. *Cognition, 212*, 104660.