Recurrent neural network and cognitive models extract different information from task behavior for predicting psychiatric traits

Wei Chen (chen-wei@g.ecc.u-tokyo.ac.jp)
The University of Tokyo
Bunkyo-ku, Tokyo 113-0033, Japan

Samuel Zorowitz (szorowi1@gmail.com)
National Institute of Mental Health
Bathesda, MD 20892, United States

Yuko Yotsumoto (cyuko@g.ecc.u-tokyo.ac.jp)
The University of Tokyo
Bunkyo-ku, Tokyo 113-0033, Japan

Stephen Keeley (skeeley1@fordham.edu)
Fordham University
New York City, NY 10458, United States

Nathaniel Daw (ndaw@princeton.edu)
Princeton University
Princeton, NJ 08544, United States

Ming Bo Cai (mingbo.cai@miami.edu)
University of Miami
Miami, FL 33146, United States

Abstract

Computational cognitive models that explicitly formulate decision-making processes with canonical learning and decision algorithms enjoy the advantage of the explainability of the estimated parameters. However, these models are highly constrained by a limited selection of model components, and fail to explain much of the variability in human decision-making behavior. Recently, deep neural networks have been applied to better capture these nuanced patterns in human decisions, utilizing their flexibility in approximating unknown data distributions. Here, we investigated the abilities of a traditional reinforcement learning (RL) model and recurrent neural networks (RNNs) to extract information from subjects' behavior in a sequential decision-making task to predict compulsivity. We found that RNN models outperform traditional RL models and feed forward DNNs. Further, to predict a static psychiatric trait from the dynamic sequential RNN states, we propose a series of novel training approaches that integrate the hidden unit activity of the RNN across trials and demonstrate that this integration choice is crucial to achieve good predictive power. Lastly, while using RNNs to directly process sequential decision-making data outperforms traditional RL parameters in predicting psychiatric traits, these two approaches may still extract different types of information from choice behavior. To test this hypothesis, we combine our RNN approach with traditional RL parameters fit from the choice data. We find that a model which combines the unconstrained RNNs trained on raw behavioral data with RL-theory extracted parameters achieves the greatest predictive power, suggesting domain-informed RL approaches are able to extract information that standard deep learning models cannot.

Keywords: recurrent neural network; reinforcement learning; computational psychiatry

Introduction

Reinforcement learning (RL) models have provided many insights into the nature of human and animal decision-making in recent decades, both for their explaining power for reward-driven learning and the neurobiological parallels of reward prediction error and expected reward values in the brain (Poeppel & Assaneo, 2020). Within RL, the two primary learning strategies are known as model-based and model-free learning (Dayan & Berridge, 2014). Model-based strategies are characterized by their use of an internal representation of the transitional probability (the model) among states of the environment in order to evaluate the value of each action. In contrast, model-free learning operates without such an internal representation, relying directly on the expected reward estimated as a consequence of any action. The two-step task is a widely adopted experimental paradigm capable of dissociating the two processes in subjects (Daw, Gershman, Seymour, Dayan, & Dolan, 2011). It involves a sequence of two-stage decisions, where in each trial each of the choice

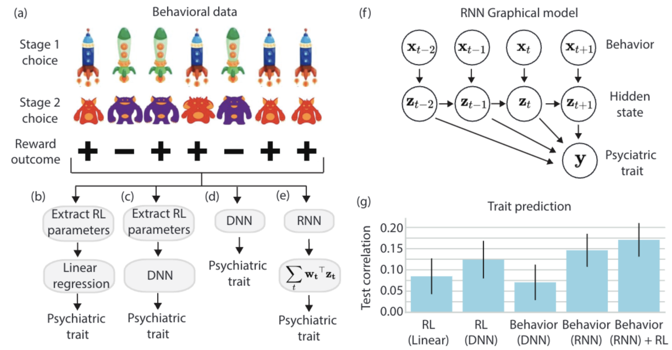

Figure 1: (a) schematic of a sequence of choices and rewards for subjects in the two-stage decision making task. (b)-(e) schematic for hypothesized models that can extract a single scalar trait value from a sequence of behaviors. (f) graphical model of the proposed RNN model for predicting psyciatric trait from sequential hidden states. (g) performance of traditional RL models, DNNs, RNNs, and RNN + RL models for predicting psyciatric symptoms from two-stage task data.

available at the first stage leads to one of two states with different probabilities, which provide different sets of available choices in the second stage. This setup allows researchers to observe whether participants adapt their first-stage decisions purely based on the final rewards regardless of the transition to the second-stage states (model-free) or if they consider the most likely states that they will reach following each first-stage choice (model-based).

Studies have revealed that reduced goal-directed control, akin to model-based learning, is associated with various mental health conditions, especially to obsessive-compulsive behavior (Gillan, Kosinski, Whelan, Phelps, & Daw, 2016). This reduction is associated with an increased reliance on habitual or model-free processes. However, previous work that showed this association has relied on a variety of simplistic constrained linear models, and thus may not characterize much of the rich variability seen in the sequential decision-making data. Here, we overcome this limitation by using recurrent neural networks to predict subjects' level of compulsivity (measured by a self-report questionnaire) based on their behavioral data in the two-step task, which allows for extraction of more complex behavioral patterns in the learning process.

Results

Our RNN model uses a vector of behavioral data from subjects (n = 975) on each trial of the task, $\mathbf{x}_t$, as input into an long short-term memory (LSTM) network to evaluate a sequential hidden state ($\mathbf{z}_t = f_\phi(\mathbf{x}_{t-1}, \mathbf{z}_{t-1})$ where ϕ represents the parameters of the LSTM). To generate a static scalar prediction for compulsivity from these hidden states, we use a linear combination of the hidden unit activity multiplied with a set of weights at all time-points to generate a prediction, $\hat{\mathbf{y}} = \mathbf{w}^\top \frac{1}{T} \sum_t a_t \mathbf{z}_t$. A schematic for our apporach can be seen in Figure 1f. We compare our RNN-based approach with four major model variations (Figure 1b-e):

1. Linear regression applied to a canonical model-free and model-based RL model parameters extracted from the behavioral data to predict psyciatric trait, similar to previous approaches in the literature (Ito & Doya, 2009; Park et al., 2017).
2. Using a feedfoward deep neural network (DNN) from these RL parameters for prediction.
3. Using a feedfoward deep neural network (DNN) from the raw trial-wise behavioral data.
4. Using an RNN on trial-wise behavioral data and a linear weighting scheme on the RNN's output.

Our findings (Figure 1g) indicate that the RNN-based models on behavioral data outperform linear or nonlinear prediction based on parameters estimated by the RL model (cross-validated). We further explored using both RNN on trial-wise behavior and RL parameters abstracted from the behavioral data to predict compulsivity symptom. Interestingly, this hybrid approach outperforms either using RNN or RL-parameters alone, even though the RNN has access to all behavioral data. This suggests that RNN and RL models may extract different aspects information that are complementary for predicting psychiatric symptoms.

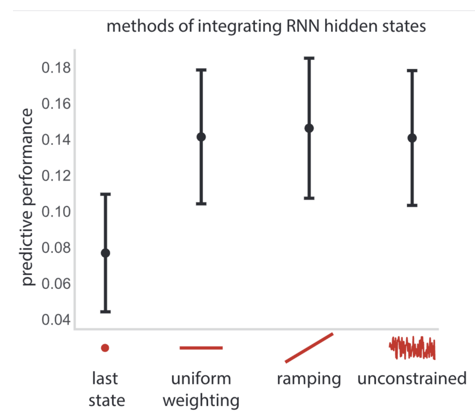

Figure 2: The performance (correlation) of using RNN on trial-wise behavioral data to predict compulsivity symptom score across different weighting schemes on RNN output.

Because RNNs even with gating mechanism may still face challenge in maintaining information early in a time series while the behavioral data are of long sequences (200 trials), we investigated different variations of our method for integrating the hidden state output $\mathbf{z}_t$ of the network across time to generate $\hat{\mathbf{y}}$. Specifically, we consider four approaches:

1. using the last hidden state (indexed by trial T) for prediction ($\hat{\mathbf{y}} = \mathbf{w}^\top \mathbf{z}_T$), the standard approach
2. using a uniform weighting across trials ($\hat{\mathbf{y}} = \mathbf{w}^\top \frac{1}{T} \sum_t \mathbf{z}_t$)
3. using a linear ramping function to scale the weights across trials ($\hat{\mathbf{y}} = \mathbf{w}^\top \sum_t (a(1 - \frac{t}{T}) + b\frac{t}{T})\mathbf{z}_t$)

4. Freely varying the weights at each trial ($\hat{\mathbf{y}} = \mathbf{w}^\top \frac{1}{T} \sum_t a_t \mathbf{z}_t$)

Our analysis show a initial evidence that the ramping approach acheives best performance in terms of testing dataset correlation (Figure 2). The underperformance of the standard LSTM RNN model, which relies solely on the last output, can potentially be attributed to the loss of important information from earlier trials in the sequence and inability of the model to retain information across long timescales. The success of the ramping method, with its fitted weights favoring earlier states (starting at 0.56 and ending at 0.21, with SEMs of 0.005 and 0.006, respectively), reinforces this perspective. The findings indicate that a nuanced approach to weighting, which considers the temporal dynamics of decision-making, is crucial in leveraging RNNs for symptom prediction, and may generalize to other contexts where sequential behavioral data are used to predict traits of subjects.

Conclusion

Here, we demonstrate that conventional deep learning methods applied to the two-stage decision making task better generates trait predictions than classical approaches. Additional work will focus on ablation of parts the decision making data to dissociate model-free and model-based strategies, as well as further exploration of how RNN and traditional RL approaches each leverage different facets of decision making data to generate psyciatric trait predictions.

References

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*(6), 1204–1215.

Dayan, P., & Berridge, K. C. (2014). Model-based and model-free pavlovian reward learning: revaluation, revision, and revelation. *Cognitive, Affective, & Behavioral Neuroscience*, *14*, 473–492.

Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *elife*, *5*, e11305.

Ito, M., & Doya, K. (2009). Validation of decision-making models and analysis of decision variables in the rat basal ganglia. *Journal of Neuroscience*, *29*(31), 9861–9874.

Park, I. A., Hwang, S.-H., Song, I. H., Heo, S.-H., Kim, Y.-A., Bang, W. S., ... Lee, H. J. (2017). Expression of the mhc class ii in triple-negative breast cancer is associated with tumor-infiltrating lymphocytes and interferon signaling. *PLoS One*, *12*(8), e0182786.

Poeppel, D., & Assaneo, M. F. (2020). Speech rhythms and their neural foundations. *Nature reviews neuroscience*, *21*(6), 322–334.