Using Recurrent Neural Networks to Understand Human Reward Learning

Mingyu Song[1], Yael Niv[1,2], Ming Bo Cai[3]

[1]Princeton Neuroscience Institute and [2]Department of Psychology, Princeton University

[3]International Research Center for Neurointelligence, The University of Tokyo

Motivation

In cognitive modeling, questions arise after model comparison:

- Have we found the best possible model?
- Can we improve the best cog model, and how?

Proposal: use recurrent neural networks (RNN) to fit behavior

- RNN serves as benchmark for goodness-of-fit
- Network analysis informs where to improve cog model

Reward learning task

(n = 102 from Mturk; Song et al., CogSci 2020)

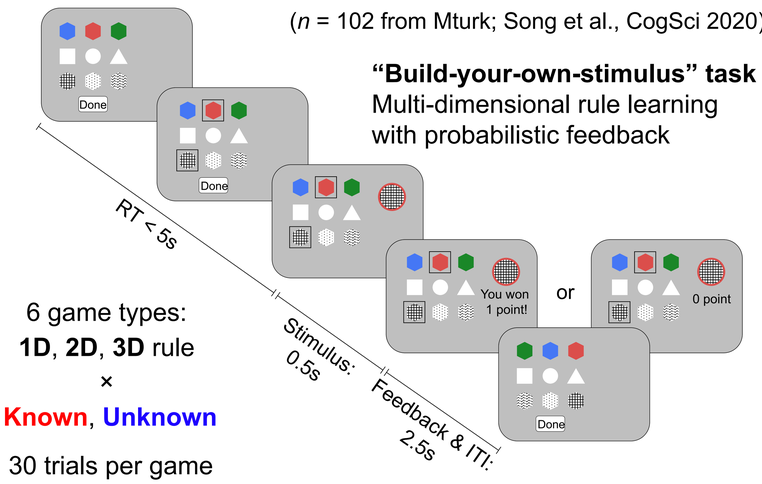

"Build-your-own-stimulus" task
Multi-dimensional rule learning with probabilistic feedback

6 game types:
1D, **2D**, **3D** rule
×
Known, **Unknown**

30 trials per game

Apply RNN to fit behavior

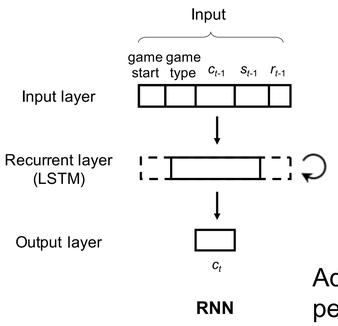

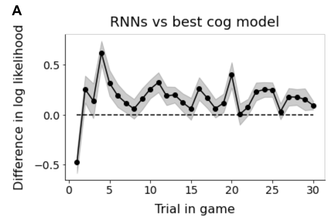

Advantage of RNN over cognitive model persisted through learning (the entire game).

RNN is better than cognitive model on predicting

- stay vs switch **(temporal dependency between trials)**
- the number of features, n_{selected} (identifying the **correct sub-space in the large choice space**)

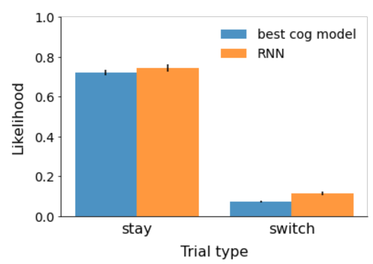

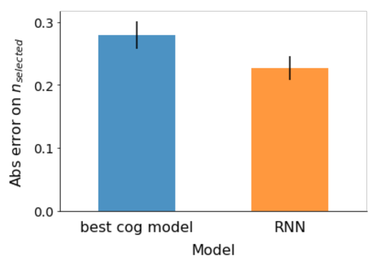

RNN embedding captures individual difference

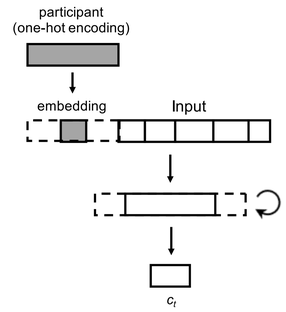

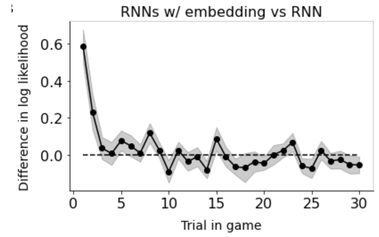

Adding embedding improves model fit, mostly in first two trials of each game.

RNN with participant embedding

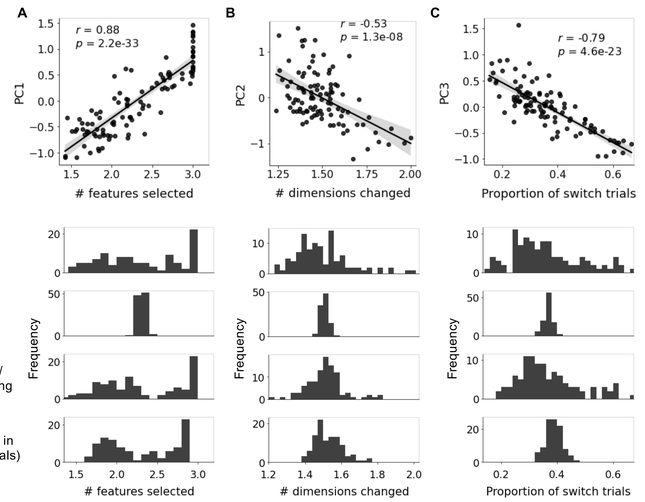

- Principle components of embedding activity correlate with cognitive variables of individual participants.
- Embedding layer necessary for individual difference.

Cognitive modeling reveals learning strategy and individual differences but fails to account for all behavioral variance

The best cognitive model

- Combines value-based (reinforcement learning) and rule-based (hypothesis testing) strategies
- Captures individual differences on hypothesis and choice policy

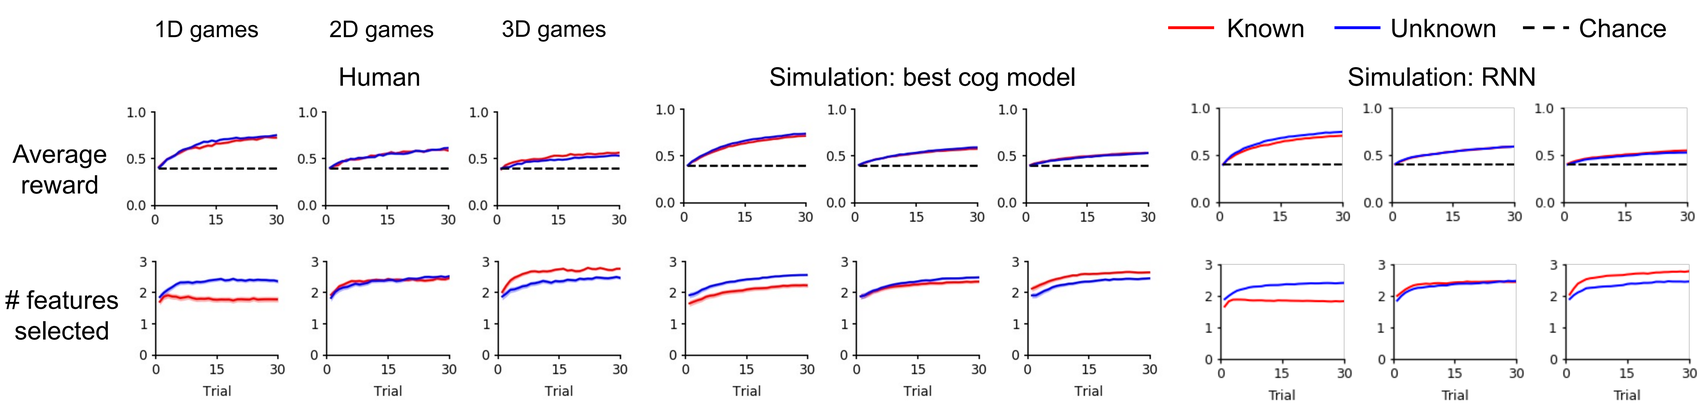

Conclusions

Use RNNs to predict human behavior in complex cognitive tasks with rich individual variability:

- Find the empirical upper bound for goodness of fit
- Reveal what is missing in the cognitive models
- Capture richness of individual behavioral differences with an embedding layer

Future work: improve cognitive model based on what we learned from RNN

Contact: mingyus@princeton.edu