Uncovering the Structure of Trial-to-Trial Variability in Perceptual Decision-Making Using Disentangled Recurrent Neural Networks

Isabelle Hoxha (isabellehoxha@gmail.com), Anne E. Urai (a.e.urai@fsw.leidenuniv.nl)
Leiden University, Wassenaarseweg 52, 2333 AK Leiden, The Netherlands

Abstract

Perceptual decision-making shows considerable trial-to-trial variability, which can be captured by latent variable models of fluctuating decision strategies. Recently, disentangled Recurrent Neural Networks (DisRNNs) have achieved data-driven discovery of such trial-to-trial latent decision strategies in multi-armed bandit tasks.

In this work, we investigate the applicability of DisRNNs for uncovering trial-to-trial structure of perceptual decision-making data. We fit DisRNNs on simulated Diffusion Decision Model (DDM) data, where the starting point or drift parameters depend on past choices. We show that the traces of the starting point and drift can be recovered in the latent variables, and that the shape of the trial-to-trial dependency of these parameters can be interpreted from the update rules learned by the network. This sets the stage for data-driven discovery of the sources of across-trial variability in real perceptual data.

Keywords: perceptual decision-making; simulations; disentangled recurrent neural network; interpretability; decision variability

Introduction

Animals must constantly make decisions about incoming sensory information: should I attack this prey from the left or the right? Do I have time to cross the street before the car arrives? Such perceptual decision-making is intrinsically variable, differing even when based on the same external information. Recently, latent variable models have been used to capture latent, time-varying decision strategies that can explain a large fraction of trial-to-trial decision variability. However, these recent models are all limited by a priori assumptions on the structure of across-trial variability. For instance, they assume deterministic update rules (Pedersen et al., 2017), discrete state switches (Ashwood et al., 2022) or slow drifts in decision parameters (Gupta & Brody, 2022)).

New interpretable recurrent neural networks (RNNs) analyze autocorrelation in the input signal to extract sparse latent variables and update rules (Miller et al., 2023), effectively allowing data-driven cognitive model discovery. In multi-armed bandit tasks, such disentangled RNNs (DisRNNs) can uncover the structure of canonical theories of reward learning (Miller et al., 2023).

These models offer an exciting opportunity for data-driven discovery of structured across-trial variability in perceptual decision-making. However, to apply DisRNNs in this new context, it should first be tested if they can jointly fit choices and reaction times (RTs) as these are crucial to disentangle different decision parameters in perceptual decision-making tasks. In this work, we thus use different variants of the Diffusion-Decision Model (DDM) to generate datasets, and assess whether DisRNNs are capable of uncovering the generative across-trial processes.

Methods

DDM simulations. We first simulated a sequence of stimuli of a perceptual decision-making task with stimuli marked as -0.2 and 0.2, to distinguish between the two choice options (e.g. leftward/rightward motion). We then simulated choices and RTs from a Diffusion-Decision Model (DDM) using PyDDM (Shinn et al., 2020).

The DDM assumes that an observer accumulates evidence from a starting point z at a drift rate ν until reaching a decision boundary $\pm B$ (here $B = 1$). This accumulation is subject to Gaussian noise (standard deviation $\sigma = 1$ here) (Figure 1).

We then incorporated structured across-trial variability in both the starting point and the drift. At each trial, the integration starting point z was updated following the equation:

$$z \leftarrow z + \alpha \times (\text{"}previous\ choice\text{"})$$

where the previous choice was set to -1 or 1, representing the lower and upper decision boundaries respectively, and the perseveration rate α was set to 0.001. z was initialized to $z_0 = 0$.

The drift rate ν was function of current stimulus and previous trial, such that:

$$\nu = d \times (\text{"}coherence\text{"}) + \beta \times (\text{"}previous\ choice\text{"})$$

where $d = 5$, and the history-dependent drift bias $\beta = 0.03$. Note that, contrary to the starting point, the drift rate equation does not accumulate past choices beyond the trial immediately preceding that one. We ran 1000 simulations of 500 trials each.

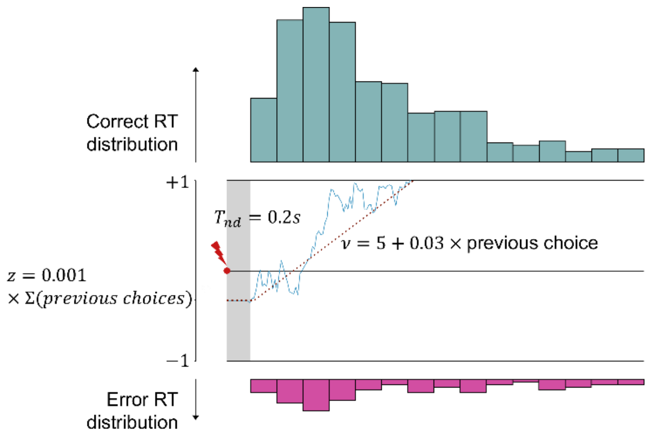

Figure 1: representation of DDM modelling and sample response time distributions from simulations.

Figure 1 summarizes the DDM choice process and displays response times for one simulation at one level of coherence.

DisRNN fitting. We extended Disentangled RNNs (Miller et al., 2023) to jointly fit choices and RTs of our simulated data (Figure 2). These networks are made interpretable by the introduction of information bottlenecks that penalize connections that use too much information from the input. That way, each latent has a simplified dependency on the inputs. To follow the structure of the dependency in the simulations, our network took the stimulus and the past choice as an input, and it had to predict the current choice and RT. We set the maximal number of latents to 5. We used mean-squared error loss to train the model. The bottleneck penalization was set to 0.001 and we ran 2,000 fitting steps. 700 simulations were used for training, and 300 for validation.

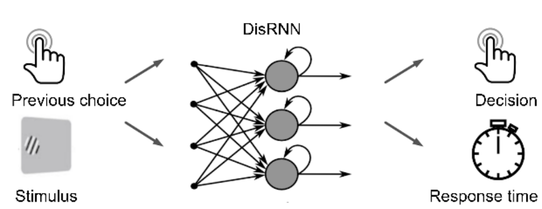

Figure 2: DisRNN inputs and outputs

Results

We fit two synthetic datasets, each with a different source of across-trial variability: starting point or drift. For each simulation, we observed that the network captured the structured variability in one latent (Figure 3A and B). Additionally, the network also learned the update rule for the parameters. Indeed, the starting point undergoes a linear update from the previous trial, by construction in the simulation, and we observed that the logit curve from the latent corresponding to the starting point is linear (Figure 3C). Similarly, the drift is constructed as an offset from previous values. This behavior is transcribed in the corresponding update rule (Figure 3D)

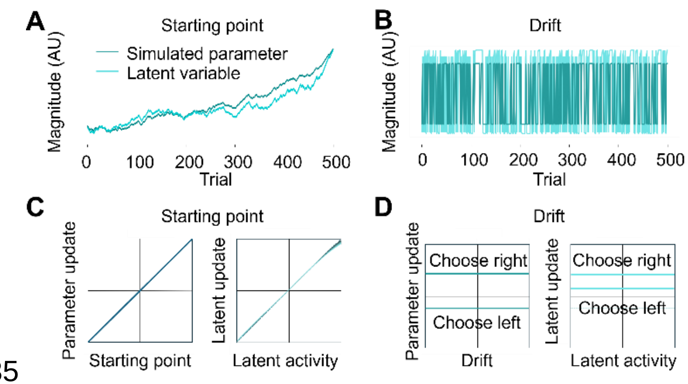

Figure 3: comparison of expressive latent evolution across trials with the starting point (A) and drift rate (B). The results are shown for one simulation. Update rules learned by the network for the starting point (C) and drift rate (D).

Future work will explore different – and more naturalistic – choice-history effects (in particular, AR processes and exponential decay of the effect of past trials with time) as well as discrete state switches (Ashwood et al., 2022).

Conclusion

In perceptual decision-making tasks, serial dependence between trials is common in behavior, despite it not being structurally imposed by the task. Here, we show that DisRNNs can uncover different generative models of serial dependence in an evidence accumulation framework: the progressive time dependence rule of the starting point and the drift parameter. The fact that the network learns the update rules of these parameters sets the for applications on real data, where sources of variability may follow yet-unknown structures.

Acknowledgements

Isabelle Hoxha is funded by the postdoctoral allocation of the Fyssen Foundation.

References

Ashwood, Z. C., Roy, N. A., Stone, I. R., The International Brain Laboratory, Urai, A. E., Churchland, A. K., Pouget, A., & Pillow, J. W. (2022). Mice alternate between discrete strategies during perceptual decision-making. *Nature Neuroscience*, *25*(2), 201–212. https://doi.org/10.1038/s41593-021-01007-z

Gupta, D., & Brody, C. D. (2022). *Limitations of a proposed correction for slow drifts in decision criterion* (arXiv:2205.10912). arXiv. https://doi.org/10.48550/arXiv.2205.10912

Miller, K. J., Eckstein, M., Botvinick, M. M., & Kurth-Nelson, Z. (2023). *Cognitive Model Discovery via Disentangled RNNs* [Preprint]. Neuroscience. https://doi.org/10.1101/2023.06.23.546250

Pedersen, M. L., Frank, M. J., & Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin & Review*, *24*(4), 1234–1251. https://doi.org/10.3758/s13423-016-1199-y

Shinn, M., Lam, N. H., & Murray, J. D. (2020). A flexible framework for simulating and fitting generalized drift-diffusion models. *eLife*, *9*, e56938. https://doi.org/10.7554/eLife.56938