

# Reliable Idiographic Parameters From Noisy Behavioral Data: The Case of Individual Differences in a Reinforcement Learning Task

Yinan Xu (yinanx97@gmail.com)

Department of Psychology, Campus Box 351525  
Seattle, WA 98195 USA

Andrea Stocco (stocco@uw.edu)

Department of Psychology & Institute for Learning and Brain Sciences, Campus Box 351525  
Seattle, WA 98195 USA

## Abstract

Behavioral data, though has been an influential index on cognitive processes, is under scrutiny for having poor reliability as a result of noise or lacking replications of reliable effects. Here, we argue that cognitive modeling can be used to enhance the test-retest reliability of the behavioral measures by recovering individual-level parameters from behavioral data. We tested this empirically with the Probabilistic Stimulus Selection (PSS) task, which is used to measure a participant's sensitivity to positive or negative reinforcement. An analysis of 400,000 simulations from an Adaptive Control of Thought - Rational (ACT-R) model of this task showed that the poor reliability of the task is due to the instability of the end-estimates: because of the way the task works, the same participants might sometimes end up having apparently opposite scores. To recover the underlying interpretable parameters and enhance reliability, we used a Bayesian Maximum A Posteriori (MAP) procedure. We were able to obtain reliable parameters across sessions (Intraclass Correlation Coefficient  $\sim 0.5$ ), and showed that this approach can further be used to provide superior measures in terms of reliability, and bring greater insights into individual differences.

**Keywords:** Probabilistic Stimulus Selection task; Reliability Test; Basal Ganglia; Direct and Indirect pathways; Computational Modeling; ACT-R

## Introduction

To understand cognition, it is important that the behavioral measures that we use to indirectly index brain function are valid and reliable. Unfortunately, this is often not the case, with published effects often showing low replicability (Bogacz et al., 2017) or task results having poor reliability across time. Idiographic (i.e., individual-level) parameters in cognitive modeling, on the other hand, can capture individual-level characteristics and are shown to have high test-retest reliability. For instance, Sense et al. (2016) have shown long-term memory rate is stable across sessions and across materials.


In this paper, we argue that idiographic parameters in cognitive modeling can be used to enhance the reliability of behavioral measures. Specifically, we show that cognitive models can be used to reliably recover the values of underlying parameters (which reflect cognitive processes) even when the behavioral data itself is noisy and lacks replicability. As an example, we will use an experimental task that has been widely adopted in neuroscience research to investigate basal ganglia function (Frank et al., 2004) but whose effectiveness has recently come under scrutiny (Baker, Stockwell, & Holroyd, 2013; Grogan et al., 2017). Specifically, we will show that the use of cognitive modeling can (a) shed light on the nature of discrepant findings by different laboratories and (b) recover interpretable, idiographic parameters from otherwise noisy behavioral data, providing superior measures of validity and reliability and greater insight into individual differences.

## The Probabilistic Stimulus Selection Task

The task examined herein is the Probabilistic Stimulus Selection (PSS) task. The PSS task is an iterative, forced-choice, implicit decision-making paradigm first introduced by Frank et al. (2004) in which participants are asked to repeatedly choose from pairs of non-verbalizable stimuli, each of which has a different probability of giving a reward (ranging from 20% to 80%). The task has a training phase and a testing phase. During the *training* phase, the participants are initially trained to select the most rewarding stimulus out of three different fixed pairs (Fig. 1, left). Feedback about the outcome of their selection (that is, whether it resulted in being rewarded or not) is shown on the screen immediately after their choice. To discourage the participants from using explicit strategies (for example by keeping a running total of each stimulus's history of successes), the stimuli are intentionally designed to be difficult to verbalize and memorize: they are represented as Hiragana characters from the Japanese writing system and are presented solely to non-Japanese speakers (for simplicity, the stimuli will be indicated with the letters *A*,

B.... F, as in Fig. 1). The learning occurring in this *training* phase is then examined in the *testing* phase, where the six stimuli are now combined into all possible pairs (Fig. 1, right) and feedback is not given to prevent further learning.

Note that, during the *training* phase, participants might learn equally well by either learning to choose the most rewarding stimuli (i.e., *A*) or by avoiding the least rewarding ones (i.e., *B*). These two processes can be distinguished in the *testing* phase by calculating two measures, Choose and Avoid accuracies. Choose accuracy is calculated as the probability of choosing *A* while paired with *C*, *D*, *E*, or *F*, and Avoid accuracy as the probability of choosing *C*, *D*, *E*, and *F* over *B*.



**Fig. 1.** Overview of the Probabilistic Stimulus Selection task. During the *training* phase (left), subjects are asked to repeatedly select one stimulus from the three possible pairs. The feedback received (“Correct!” or “Incorrect!”) depends on the stimulus chosen and is shown immediately after each choice. The six stimuli are presented in fixed pairings. During the *testing* phase (right), subjects perform the same task as the training phase but without the feedback. The stimuli now appear in new pairings that include either the most rewarding stimulus (green lines) or the least rewarding stimulus (red lines) against each of the remaining stimuli.

The importance of this task lies in the fact that Choose and Avoid accuracies provide insight into a person’s biology, and, specifically, into the physiology of the basal ganglia. The basal ganglia are a set of subcortical nuclei that modulate the activity of the prefrontal cortex and are involved in many cognitive functions, most importantly in acquiring procedural knowledge (Knowlton & Squire, 1994). The connections between these nuclei are organized into two pathways, called the direct and the indirect pathway, which have opposite effects on cortical activity (Albin et al., 1989; DeLong, 1990). While the direct pathway exerts an excitatory effect on the prefrontal cortex, the indirect pathway has an inhibitory influence. The striatal neurons that originate the two pathways also express different dopamine receptors: While the direct pathway neurons express d1 receptors that are excited by dopamine release, indirect pathway neurons expressed d2 receptors and are inhibited by dopamine (Gerfen et al., 1990). Because dopamine is important in reward-based learning and decision-making, it was hypothesized that Choose

accuracy reflects the contribution of the direct pathway and Avoid accuracy reflects the activity on the indirect pathway. In fact, the study by Frank et al. (2004) with Parkinson’s Disease (PD) patients shows that, when on Dopamine-promoting medication, patients are more likely to be “Choosers”, meaning their Choose accuracy is higher than the Avoid accuracy. Correspondingly, they are more likely to be “Avoiders” when off medication and their dopamine level is low. Evidence from Frank et al. (2007) also shows that higher Choose accuracy is associated with people that have DARPP-32 gene polymorphisms that promote the expression of d1 Dopamine receptors on the direct pathways, and higher Avoid accuracy is associated DRD2 gene polymorphisms that promote expression of d2 Dopamine receptors on the indirect pathways.

However, the PSS task’s reliability has recently been called into question. Experiments by Grogan et al. (2017) with PD patients failed to reproduce effects of dopaminergic medications on PSS performance, in contrast to previous studies done by Frank et al. (2004; 2007). Also, Baker, Stockwell, and Holroyd (2013) found no evidence that patterns of behavior are stable in this task over time. In their study, they conducted a test-retest reliability analysis on the PSS task performance on 90 undergraduate students. This result showed poor reliability of the behavior measures in indexing cognitive processes in reinforcement learning.

## Summary

In summary, although existing literature suggests that the PSS task can successfully track the function of the basal ganglia’s direct and indirect pathways, and the task has been therefore vastly used for this purpose, the reliability of the task needs to be further determined. To deal with this matter, we conducted a new reliability test (Experiment 1) on the same versions of the PSS task used in Frank, Seeberger, and O’Reilly, (2004) and used a computational model of this task to examine why poor reliability exists (Experiment 2). Furthermore, we aimed to recover important individual differences information from behavioral data using this model-based approach. In this case, a Maximum A Posteriori (MAP) procedure was done to recover the underlying parameters of the behavioral data.

## Behavioral Experiment

### Method and Materials

**Participants** 71 healthy participants (age 18-30, 41 females) from the University of Washington’s undergraduate population took part in the experiment in exchange for credit course. All participants completed two sessions of the PSS exactly one week apart. The second session always occurred on the same day of week and at the same time of day as the first session.

**Task** All participants completed the PSS task in the same version used by Frank et al. (2004). Participants were asked to place their left index finger on button “1” and right index

finger on button “0” of a standard computer keyboard placed in front of them. They were noticed about the procedure. Pairs of Hiragana characters were then shown on the screen with a fixation point in between each trial. They then pressed the button corresponding with the characters they intuitively think would be correct. Feedbacks were shown on the screen after each selection during the training phase as “Correct!” in blue color or “Incorrect!” in red color. If the participant didn’t press any button within 6 seconds, “no response detected” in red color was shown on the screen. This was to ensure that the subject was engaging in the task and also to discourage the subject from using explicit methods to remember the patterns rather than learning them through trial and error. After a maximum of six repetitions of the training phase, participants moved to the testing phase where their Choose and Avoid accuracies were measured.

## Results

**Split-test Reliability** First, we examined the split-test reliability of these measures. This was done by separately calculating the values of the two main variables, Choose and Avoid, for different pairs of stimuli, depending on whether *A* and *B*, were presented on the left (e.g., “AC”, “BC”, etc.) or on the right (e.g., “CA”, “CB”, etc.). These measures were called Choose Left, Avoid Left, Choose Right, and Avoid Right, respectively. The Pearson correlation between the Left and Right version of each measure was calculated. As shown in Fig. 2, the split-test correlation coefficients of Choose and Avoid were significant in both sessions. Specifically, we found a positive correlation of Choose Left and Choose Right in Session 1 [ $r(71) = .44, p < .001$ ] and Session 2 [ $r(71) = .40, p < .001$ ], and between Avoid Left and Avoid Right in Session 1 [ $r(71) = .46, p < .001$ ] and in Session 2 [ $r(71) = .40, p < .001$ ].


**Test-retest Reliability** Then, we examined the test-retest reliability across sessions of the same measures. In contrast to the split-test correlations, no significant correlation was found for either Choose [ $r(71) < 0.10, p > 0.60$ ] or Avoid [ $r(71) = 0.15, p > 0.20$ ] across sessions (Fig. 3).

**Intraclass Correlation Coefficient** Finally, for each of the measures of interest, we also calculated the intraclass correlation coefficient (ICC: Shrout and Fleiss, 1979). ICC measures the proportion of variance in the measures of interest against the total variance, and is used as an assessment of the consistency of quantitative measurements between sessions. In our study, the variance of interest is between different measurements *M* of the same variables (Choose or Avoid) across left/right presentation or sessions, and the total variance is due to both *M* and the individual participants *P*. Thus:


$$ICC = \sigma_M^2 / (\sigma_M^2 + \sigma_P^2)$$

As shown in Fig. 3, although the ICC values for Choose and Avoid between sessions were greater than zero, they were also markedly inferior to their split-half counterparts

and both values fell below the 0.40 threshold indicated by Cicchetti (1994) as “poor” reliability.



**Fig. 2.** Split-test and Test-retest reliability of Choose and Avoid accuracies. Red lines represent significant, and black lines represent non-significant, correlations.



**Fig. 3.** ICC results. Choose (red) and Avoid (blue) over sessions (Right column) showed poor consistency.

## Summary

An analysis of the Choose and Avoid measures in the PSS task has yielded somewhat contrasting results. Across sessions, both measures show very poor reliability as indexed by both Pearson correlations and ICC values. This finding is in line with the reports of Bogacz (2017) and Holroyd (2013), which called into question the original results by Frank. On the other hand, the same measures had good reliability scores *within* each session ( $r > 0.4$ ,  $ICC \geq 0.5$ ), suggesting that the two measures were *not* intrinsically unreliable.

Two possible explanations exist for these findings. One is that Choose and Avoid do not index any underlying stable feature of a participant's biology (such as the relative strengths of their basal ganglia pathways) but some other characteristic that is reliable only within a single session. This could be, for example, mental states such as fatigue. Another hypothesis is that Choose and Avoid might be intrinsically noisy indicators of the underlying basal ganglia activity. For instance, for a participant with high learning rates for both pathways, Choose might dominate in one session while Avoid might dominate in another.

To distinguish between these two hypotheses, we examined the performance of an existing model of the task and we applied Bayesian methods to estimate the most likely underlying model parameters for each participant.

## Computational Model


### Methods and Materials

**Computational Model** A model of the PSS was recently published (Stocco, 2018) and its code made available online<sup>1</sup>. The model was developed using the Adaptive Control of Thought - Rational (ACT-R) architecture, which is currently the most common cognitive architecture in use (Kotseruba & Tsotsos, 2018). Similar to other architectures, it contains vector-like structures called “chunks”, which are used to represent static information like semantic memory (“a dog is walking”), visual input (“red rectangle on the left”), or motor commands (“press the green button”). These chunks are then placed into specialized modules (such as “vision”) where they become accessible to procedural knowledge (represented as “production rules” or “productions”) to carry out cognitive and motor actions. Only one production is selected at any given time, only one production is allowed to fire; this production is selected amongst competing rules on the basis of its relative *utility*, a scalar value that represents the estimated future rewards generated by their applications and is learned through a reinforcement learning algorithm.

The model by Stocco (2018) assumes that task performance relies entirely on procedural knowledge. This is due to the reinforcement learning nature of the task, and the mapping between procedural knowledge in ACT-R and

Basal Ganglia (Anderson, 2005; Anderson et al., 2008; Stocco & Anderson, 2008). In addition, and as previously noted, the PSS task is designed to exclude the possibility that participants were relying on declarative knowledge. In fact, experimental results show task performance is invariably affected by manipulation of basal ganglia function and not affected by manipulation on the formation of declarative memories (Frank, O'Reilly, & Curran, 2006). Therefore, this ACT-R model implemented a procedural-only approach for the PSS task. The model also captured competitive action selection, which also corresponds to the basal ganglia's biological computations. Once the stimuli on the screen are encoded as a visual chunk, all the productions that match the current stimuli compete for execution. To capture the competition between the basal ganglia's direct and indirect pathways, the model also uses opposing and competing “Choose” and “Avoid” productions for each stimulus (Fig. 4). For example, if *A* is presented, both “Choose *A*” and “Avoid *A*” would compete to select either stimulus *A* or the other stimulus on the screen (*B* in Fig. 4).

The functioning of the model is governed by four parameters only:  $\alpha$ ,  $s$ ,  $D1$ , and  $D2$  (Fig. 4). The model parameter  $\alpha$  represents the learning rate in reinforcement learning, that is, how much each production's utility is adjusted for after each feedback. The parameter  $s$  represents the noise in selecting each action at the decision phase. Finally,  $D1$  and  $D2$  model the density of dopamine  $d_1$  and  $d_2$  receptors in the basal ganglia's direct and indirect pathways, respectively, and modulate the effects of the learning rate for the Choose and Avoid production rules. Thus,  $D1$  and  $D2$  represent the unobserved quantities that the Choose and Avoid measures are purported to operationalize.



**Fig. 4.** Overview of the PSS model performing a sample trial of the PSS task.

<sup>1</sup> [http://github.com/UWCCDL/PSS\\_model](http://github.com/UWCCDL/PSS_model)



**Simulations** The original paper (Stocco, 2018) provides values for the learning rate  $\alpha$  and noise parameter  $s$  for the general population, as well as the distribution of values of D1 and D2 that capture the observed variability in healthy individuals. In this study, we used the values of  $\alpha = 0.018$  and  $s = 0.1$  (which were fit to the healthy control data in Stocco, 2018), and parametrically varied the values of D1 and D2 from 0 to 2 in increments of 0.05 (which were used in Stocco, 2018, to capture individual differences in the PSS task). For each combination of D1 and D2 parameter values, the model was then run 250 times, and the probability distributions of each combination of Choose and Avoid measures were recorded. A total of  $41 \times 41 \times 250 = 420,250$  simulations were run.

#### Maximum A Posteriori (MAP) Parameter Estimation

The simulations described in the previous section provide an estimate of the likelihood of observing a particular behavioral outcome  $Y$  (that is, a combination of Choose and Avoid values) given  $\theta$  (that is, a combination of values for the D1 and D2 parameters). In addition to these likelihood estimates, we were interested to explore whether the model's simulations could be used to estimate reliable values for D1 and D2 (the unobservable parameters that govern learning rates in the two pathways) from the observable behavioral measures (Choose and Avoid). To do so, we fitted the model to each individual participant using a Bayesian Maximum A Posteriori (MAP) procedure. In Bayesian statistics, a MAP is defined as the estimate of the maximum likelihood of an unobservable quantity on the basis of both the empirically observed data and a prior hypothesis about the distribution of that quantity. In our case, the procedure was used to recover the most likely values of  $\theta$  (D1 and D2) given the observed Choose and Avoid values of a given participants ( $Y$ ), that is,  $\text{argmax}_{\theta} P(\theta | Y)$ . Using Bayes theorem, this quantity can be rewritten as:


$$P(\theta | Y) = \text{argmax}_{\theta} [P(Y | \theta) \times P(\theta) / P(Y)].$$

The likelihood values  $P(Y | \theta)$ , that is, the distributions of Choose and Avoid accuracies given pairs of D1 and D2 parameter, can be directly computed from the model simulations (Fig. 5). To estimate the parameter priors  $P(\theta)$ , we followed the following logic. First, we modeled the probability distribution of each parameter value as a normal distribution  $N(\mu, \sigma)$  with mean  $\mu = 1$  and  $\sigma = 0.5$ . This captures the finding that the values for D1 and D2 that best represent variability among healthy participants vary between 0.5 and 1.5, with 1 being the population mean (Stocco, 2018). The joint probability distribution of D1 and D2 was then modeled by setting the correlation between the two distributions to  $r = 0.5$ . The correlation between these two parameters is suggested by the facts that, in the basal ganglia, the distribution of receptors is not independent, and that the activity in both receptors is driven by a common force (the release of dopamine), and that, the recorded activity of the two pathways is anticorrelated.

With both the likelihood and the joint probability distribution in place, we calculated the MAP estimates for D1 and D2 for each participant.


## Results

**Likelihood Distributions** An inspection of likelihood distributions the model's performance provided a first insight into the reasons for the poor test-retest reliability of the PSS task measures. Under certain combinations of parameters, the model tends to converge on the same estimates of Choose and Avoid; this result is represented by a likelihood distribution with a unique global maximum (Fig 5, left). However, under *most* combinations, the likelihood distribution did *not* have a single maximum, and the model would move towards different values for Choose and Avoid in different runs. (Note that, since there is a finite number of pairs in the *testing* phase, the values of Choose and Avoid are discrete, and the multiple peaks in Fig. 5 do not represent an approximation due to the discretization of continuous variables). This suggests that the poor reliability of Choose and Avoid accuracies might be due to the nature of the task and the ways the measures are calculated.




**Fig. 5:** Variability in the likelihood of possible results  $Y$  as a function of different parameter values  $\theta$ . When  $D1 = D2 = 0.5$  (Left), the model converges on a single global maximum (Choose = 0.6125, Avoid = 0.6125). However, when  $D1 = D2 = 2.0$  (Right), multiple possible results (i.e. maxima) are equally likely. Colors represent probability densities for each Choose/Avoid accuracy combination.

**Test-Retest Reliability** After calculating the MAP parameter estimates for each participant, we applied the same test-retest reliability analyses that were used for the behavioral measures to the individual parameter values. In contrast to our behavioral findings (Fig. 2), we found statistically significant Pearson correlations across sessions for both D1 [ $r(71) = 0.33, p < 0.005$ ] and D2 [ $r(71) = 0.35, p < 0.003$ ] (Fig. 6). Similar results were found for the corresponding ICC values, with the values for D1 (ICC=0.49) and D2 (ICC=0.51) being more than twice as large as the corresponding values for Choose and Avoid, and within Cicchetti's (1994) range of "fair" reliability (Fig. 7).



**Fig. 6.** Correlation between MAP estimates of the underlying *D1* and *D2* model parameters across sessions for all participants. Red lines indicate significant correlations.



**Fig. 7.** A comparison of the Intraclass Correlation Coefficient (ICC) values of the behavioral measures (left) and the MAP estimates of the underlying model parameters (right) that index the function of the two pathways.

## Conclusion

This paper has presented evidence that computational cognitive models can be used to recover interpretable and reliable parameters from noisy behavioral data, making up for discrepant findings. More specifically, the study examined the reliability of two commonly used measures in the PSS task and suggested recovering the underlying parameters by fitting cognitive models can be used in the future as an alternative to behavioral data analysis.

The study consisted of two parts. First, we tested the reliability of the PSS task's most important measures, Choose and Avoid accuracies, and showed that their test-retest reliability is poor. Across two sessions one week apart, the Choose and Avoid values were uncorrelated within participants. Second, we adopted a computational modeling approach to gather a better understanding of the data. An ACT-R cognitive architecture model was then implemented here. The model simulated the competitive dynamics of the two basal ganglia pathways that together drives the reinforcement learning process. A set of parameters generated by the model was then manipulated to generate probability distributions for every possible Choose

and Avoid value. With evidence-based idea of how *D1* and *D2* distributed, when fit the actual data obtained from participants in the first study back to these probability distributions using Maximum A Posterior, we got the joint probability distribution for *D1* and *D2* of our participants. Recalculating test-retest reliability and ICC with these parameters showed great improvement. As a result, both *D1* and *D2* show correlation across sessions, meaning a better reliability than Choose and Avoid measures.

With the PSS task as a starting point, we see this proposed approach as an exciting approach in future studies as well. First and foremost, using cognitive models suggests a new and more reliable way of dealing with behavioral data, as idiographic parameters reflect stable individual traits that possess high test-retest reliability. It also highlights, when understanding cognition, the importance of looking for underlying idiographic parametric indices for a more direct reflection on cognitive processes. Idiographic parameters can be used to generalize or predict behaviors across tasks, for instance, Lovett and colleagues were able to estimate attentional spreading activation from a working memory task, and use it to predict performance in a second task (Daily et al., 2001; Lovett et al., 2000). This approach can also be used to reflect individual's measurable biological features, such as procedural learning rate and the density of dopamine receptors (Stocco 2018; Stocco et al., 2017), providing a platform for further studying individual differences, and offering potential benefits toward the goal of predicting individual performance and its development across times.

## References

- Albin, R. L., Young, A. B., & Penney, J. B. (1989). The functional anatomy of basal ganglia disorders.
- Anderson, J. R. (2005). Human symbol manipulation within an integrated cognitive architecture. *Cognitive Science*, 29(3), 313–341.
- Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Sciences*, 12(4), 136–143.
- Anderson, J. R., Taatgen, N. A., & Byrne, M. D. (2005). Learning to achieve perfect timesharing: Architectural implications of Hazeltine, Teague, and Ivry (2002). *Journal of Experimental Psychology: Human Perception and Performance*, 31(4), 749–761.
- Baker, T. E., Stockwell, T., & Holroyd, C. B. (2013). Constraints on decision making: implications from genetics, personality, and addiction. *Cognitive, Affective, & Behavioral Neuroscience*, 13(3), 417–436.
- Cicchetti, D. V. (1994). Guidelines, criteria, and rules of thumb for evaluating normed and standardized assessment instruments in psychology. *Psychological assessment*, 6(4), 284.
- 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.
- DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in neurosciences*, 13(7), 281-285.
- Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T., & Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proceedings of the National Academy of Sciences of the United States of America*, 104(41), 16311–16316.
- Frank, M. J., Samanta, J., Moustafa, A. A., & Sherman, S. J. (2007). Hold your horses: Impulsivity, deep brain stimulation, and medication in parkinsonism. *Science*, 318(5854), 1309–1312.
- Frank, M. J., Seeberger, L. C., & O'reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. *Science*, 306(5703), 1940–1943.
- Frank, M. J., O'Reilly, R. C., & Curran, T. (2006). When memory fails, intuition reigns: Midazolam enhances implicit inference in humans. *Psychological Science*, 17(8), 700–707.
- Gerfen, C. R., Engber, T. M., Mahan, L. C., Susel, Z. V. I., Chase, T. N., Monsma, F. J., & Sibley, D. R. (1990). D1 and D2 dopamine receptor-regulated gene expression of striatonigral and striatopallidal neurons. *Science*, 250(4986), 1429-1432.
- Grogan, J. P., Tsivos, D., Smith, L., Knight, B. E., Bogacz, R., Whone, A., & Coulthard, E. J. (2017). Effects of dopamine on reinforcement learning and consolidation in Parkinson's disease. *Elife*, 6, e26801.
- Kotseruba, I., & Tsotsos, J. K. (2018). 40 years of cognitive architectures: Core cognitive abilities and practical applications. *Artificial Intelligence Review*, 1-78.
- Lovett, M. C., Daily, L. Z., & Reder, L. M. (2000). A source activation theory of working memory: Cross-task prediction of performance in ACT-R. *Cognitive Systems Research*, 1(2), 99-118.
- Schutte, I., Slagter, H. A., Collins, A. G., Frank, M. J., & Kenemans, J. L. (2017). Stimulus discriminability may bias value-based probabilistic learning. *PloS one*, 12(5).
- Sense, F., Behrens, F., Meijer, R. R., & van Rijn, H. (2016). An individual's rate of forgetting is stable over time but differs across materials. *Topics in cognitive science*, 8(1), 305-321.
- Shrout, P. E., & Fleiss, J. L. (1979). Intraclass correlations: uses in assessing rater reliability. *Psychological bulletin*, 86(2), 420.
- Stocco, A., Murray, N. L., Yamasaki, B. L., Renno, T. J., Nguyen, J., & Prat, C. S. (2017). Individual differences in the Simon effect are underpinned by differences in the competitive dynamics in the basal ganglia: An experimental verification and a computational model. *Cognition*, 164, 31-45.
- Stocco, A. (2018). A biologically-plausible action selection system for cognitive architectures: Implications of basal ganglia anatomy for learning and decision making models. *Cognitive Science*, 42(2018), 457–490.