

Configural Learning depends on Task Complexity and Temporal Structure

Nicholas Menghi (N.Menghi@uea.ac.uk)

School of Psychology, University of East Anglia, Norwich NR4 7TJ. UK

Will Penny (W.Penny@uea.ac.uk)

School of Psychology, University of East Anglia, Norwich NR4 7TJ. UK

Abstract

This paper describes a set of associative learning experiments in which the appropriate response depends on multiple relevant stimuli. We vary both the complexity of the stimulus-response mapping (task) and the temporal structure of the stimuli that are presented. We find that both of these manipulations affect the accuracy with which the task can be learnt, and that task complexity affects the proportion of subjects who correctly provide declarative knowledge of the underlying association. Computational modelling of subjects' behaviour, based on Dynamic Logistic Regression models, allowed us to probe the strategies that subjects employed during learning. We found that the majority of subjects employed a configural learning strategy during the complex task and a mixed configural/rule-based strategy during the simpler task. Computational modelling also provided an entropy-based index of strategy exploration with greater exploration observed during the complex task.

Keywords: decision making; configural learning; conjunctive representation; weather prediction task; dynamic logistic regression; rule-based

Introduction

Associative learning of stimulus-response mappings can proceed using a variety of cognitive strategies, neuronal representations, and decision making systems (Domenech & Koehlin, 2015). This paper considers a learning context in which there are multiple relevant stimuli and the appropriate response depends on the pattern of stimuli that are presented. This type of learning has previously been studied, for example, using the Weather Prediction Task (WPT) (Knowlton, Squire, & Gluck, 1994). The multiple systems that are engaged during this task include a medial temporal lobe representational system and a habitual feedback-based fronto-striatal decision system (Poldrack et al., 2001).

In a recent study Duncan, Doll, Daw, and Shohamy (2018) created two WPT-based learning tasks. The first required an 'elemental' strategy. Here, one learns the log-odds with which the constituent elements of a pattern determine the outcome. The log-odds of the outcome are then given by the sum of the log-odds of the constituent elements - the "whole being the sum of its parts". The second task required a 'configural' strategy where the log-odds of the outcome was specific to each pattern (or 'configuration'). They found that most subjects used a configural strategy when required (and that it relied on conjunctive representations in anterior hippocampus).

Interestingly, approximately half of subjects used a configural strategy when the simpler elemental strategy would have sufficed. This suggests that configural learning is automatically engaged and rather effortless.

In this paper we describe a new set of WPT-based learning tasks which are designed to probe the characteristics of this configural learning system. This work is motivated by the concern that both configural and elemental strategies are rather limited as general computational strategies. Firstly, elemental decision making makes the limiting assumption that the log-odds of the outcome is a linear function of the stimuli. A fully configural learning system on the other hand allows for mappings with arbitrarily complex non-linearities, but this won't scale-up because the number of configurations grows exponentially with number of stimuli.

We therefore designed our WPT-tasks to have true stimulus-response mappings that were nonlinear but had intermediate levels of complexity (between elemental and configural). Additionally, although the tasks were probabilistic each could be described using a (simple or complex) verbal rule (Ballard, Miller, Piantadosi, & Goodman, 2017). We hypothesized that subjects would abandon a configural strategy in favour of a rule-based strategy, especially for the simpler task. We also predicted that the rule subjects implicitly learnt would at some stage become explicit so that subjects would be able to declare the rule they were using.

The experiments in Duncan et al. (2018) used only 6 configurations. We scale this up to 25 in this paper so as to push the limits of the configural learning system. Another unrealistic feature of WPT concerns the temporal structure of the stimuli. In the standard WPT the probability that any particular pattern is presented on any given trial follows a uniform distribution. The patterns that appear to our sensory cortices, however, have a good deal of temporal regularity and we hypothesized that this could affect both the learning strategy implemented and the accuracy with which the task could be performed. We therefore manipulated the temporal structure of the stimuli in our experiments and hypothesized this would also affect declarative knowledge.

Methods

60 participants performed a WPT-based learning task in which they used feedback to learn which weather outcome was associated with what pattern of visual cues. At each trial two geometrical shapes, each having three, four, five, six or seven sides, were presented. There are thus $k = 1..25$ unique patterns. Participants had to classify the pattern with a weather



outcome, Sun or Rain, (procedure in Figure 1 panel A). We assessed whether participants had learnt the task implicitly or explicitly by asking them about their knowledge after each task had been completed.

We manipulated task difficulty (within-subjects factor), and the temporal structure of the cues (between-subjects factor). Task difficulty varied over two levels such that a "simple task" could be described using a single logical clause whereas a "complex task" could only be described using multiple logical clauses. The probabilistic structure of these two tasks was specified by making the log-odds of the outcome a quadratic function of stimulus characteristics. Flipping the sign of a single parameter in this mapping produced either the simple or complex rule as shown in Figure 1 panel B by the signed dependency structure.

Cues were presented with three different temporal structures. The first was generated as is standard in the WPT such that the probability of the k th cue pattern occurring was uniform over trials, $p(u_t = k) = 1/T$ where $T = 250$ is the number of trials per task. The second structure, was generated to create a "blocky" design such that the probability of the k th cue pattern occurring was concentrated within an interval. We used the distribution $p(u_t = k) = N(\mu_k, \tau) / \sum_j N(\mu_j, \tau)$. We refer to this as an 'Interval' structure by analogy with interval training in physical exercise. The third temporal structure was a mixture of the first two. The first 175 trials were generated from an interval distribution, the last 75 from a uniform one. We conceived of this last section of trials as a revision period in which knowledge could be created or consolidated.

The 60 subjects were split into 20 who were exposed to a uniform structure, 20 to an interval structure and 20 to a mixed structure. Each subject learnt both simple and complex tasks.

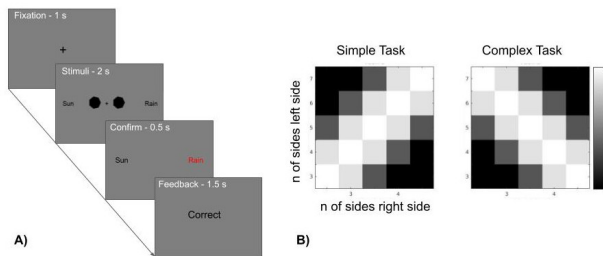


Figure 1: (A) The structure of each trial. (B) For the simple task the subject should decide "Sun" if the number of sides of the two geometrical shapes was the same. For the complex task the subject should decide "Sun" if the total number of sides was equal to ten.

Computational Model

Subjects behaviour was modelled using a Dynamic Logistic Regression (DLR) framework (Speekenbrink, Channon, & Shanks, 2008) in which the outcome on each trial, $y_t = \{1, 0\}$

for $\{Sun, Rain\}$, was modelled as

$$g_t \equiv p(y_t = 1) = \frac{1}{1 + \exp(-\lambda v_t)} \quad (1)$$

where λ is a subject-specific decision noise parameter and the activations are given by

$$v_t = w_t^T x_t \quad (2)$$

The regression coefficients were estimated online using gradient ascent as

$$w_t = w_{t-1} + \alpha x_t (y_t - g_t) \quad (3)$$

where α is a subject-specific learning rate.

A Configural model ($m = 1$) is specified by choosing x_t to be a $[K \times 1]$ binary vector with k th entry equal to 1 if the k th pattern was presented on trial t . This model therefore has 25 regression coefficients.

A Rule model ($m = 2$), is specified by choosing x_t to be a two element binary vector with entries $[1, 0]$ if condition "C" is met, otherwise $x_t = [0, 1]$. Here "C" is either "The two shapes have the same number of sides" (Task 1) or "The sum of the sides of the two shapes is 10" (Task 2). This model therefore has 2 regression coefficients.

A Mixed model ($m = 3$) is a mix between the rule and the configural model, where if the rule is met $x_t = [1; \text{zeros}(20, 1)]$ (in matlab notation). If the rule condition is not met then x_t is set to the all zeros vector but with a 1 in position $k + 1$ if pattern k was presented. This model therefore has 21 regression coefficients.

Subject specific parameters, $\theta = \{\lambda, \alpha\}$, were estimated using Bayesian inference where the log likelihood of subject decisions, $A = \{a_1, \dots, a_t, \dots, a_T\}$, is given by

$$\log p(A|\theta, m) = \sum_{t=1}^T p(a_t|\theta, m) = \sum_{t=1}^T a_t \log g_t + (1 - a_t) \log(1 - g_t) \quad (4)$$

and Gaussian priors were used over LogIt functions of parameters (Mathys, Daunizeau, Friston, & Stephan, 2011) to define prior densities, $p(\theta|m)$. A Laplace approximation was used to compute the evidence for each model, $p(A|m)$, which facilitated Bayesian Model Comparison.

We also computed a running estimate of the model probabilities as each task progressed. This assumed a flat prior over models at $t = 0$, that is $p(m_t) = 1/3$ indicating subjects had no initial preference for the configural, rule or mixed models, which was then updated recursively using Bayes rule

$$p(m_t = i) = \frac{p(a_t|m_{t-1} = i)p(m_{t-1} = i)}{\sum_j p(a_t|m_{t-1} = j)p(m_{t-1} = j)} \quad (5)$$

Finally, we computed the entropy over this distribution

$$h_t = \sum_i p(m_t = i) \log_2 \left(\frac{1}{p(m_t = i)} \right) \quad (6)$$

$$H = \sum_t h_t$$

where h_t is a trial-by-trial measure and H is the average entropy over a task which we use as an index of strategy exploration.

Results

Declarative Knowledge

For the simple task 28 out of 60 subjects declared the correct strategy whereas only 1 out of 60 did for the complex task. This significant difference ($\chi^2 = 33.14, p < 0.001$) confirms that our difficulty manipulation was successful. There was no significant variation in number of declarations as a function of Temporal Structure - see Table 1.

	Uniform	Interval	Mixed
Simple	10	8	10
Complex	0	0	1

Table 1: **Knowledge** Number of subjects who explicitly declared the correct strategy by task and temporal structure.

Overall Accuracy

We define overall accuracy as the correct rate averaged over all 250 trials. We ran a two-way mixed design ANOVA with dependent variable overall accuracy and factors of (i) difficulty and (ii) temporal structure. This revealed both main effects, interaction not significant (Difficulty, $F(1) = 21.576, p < 0.001$; Temporal structure, $F(2) = 3.908, p = 0.025$). See Table@2 for full results.

Participants performed better in the simple task compared to the difficult one (see Figure 2 A). Also, participants performed differently based on the temporal structure they were presented with (see Figure 2 B). These significant differences show that our manipulations were successful. Tests of simple effects showed that accuracy was higher in the interval than uniform condition for the complex task ($t(19) = 3.076, p = 0.006$) but not for the simple task. Participants who explicitly declared the strategy performed better ($t(58) = 2.282, p = 0.026$).

Inferred Strategy

Bayesian model comparison revealed that during the simple task, most of the participants used the mixed strategy (34/60), only 10/60 used the configural strategy and 16/60 the rule strategy. Conversely, during the complex task 48 participants used the configural strategy, 11 the rule strategy and just a single participant the mixed strategy. Tables 3 and 4 give a

	Uniform	Interval	Mixed
Simple	61.04(1.76)	64.07(1.33)	63.56(1.59)
Complex	54.92(1.49)	61.14(1.68)	57.33(1.08)

Table 2: **Accuracy** Mean accuracy per task difficulty and temporal structure. Standard error of the mean between brackets.

breakdown of the inferred strategy as a function of temporal structure.

	Uniform	Interval	Mixed	Tot
Configural ($m = 1$)	2	6	2	10
Rule ($m = 2$)	4	5	7	16
Mixed ($m = 3$)	14	9	11	34

Table 3: **Simple Task** Number of subjects favouring each strategy as a function of temporal structure. Model/Strategy m was deemed favoured by a subject if it had the highest Model Evidence, $p(A|m)$.

	Uniform	Interval	Mixed	Tot
Configural ($m = 1$)	18	17	13	48
Rule ($m = 2$)	2	3	6	11
Mixed ($m = 3$)	0	0	1	1

Table 4: **Complex Task** Number of subjects favouring each strategy as a function of temporal structure. Model/Strategy m was deemed favoured by a subject if it had the highest Model Evidence, $p(A|m)$.

Strategy Exploration

A two-way mixed ANOVA with temporal structure as a between-subjects factor and task difficulty as a within-subjects factor revealed both main effects, but no interaction (Temporal structure, $F(2) = 11.422, p < 0.001$; Task Difficulty, $F(1) = 4.239, p = 0.044$). Entropy, H , is higher in the complex task, indicating more exploration, as compared to the simple task. Furthermore, entropy is higher with uniform temporal structure as compared to both the interval ($t(38) = 4.777, p < 0.001$) and the mixed structures ($t(38) = 2.337, p = 0.024$), entropy is higher with mixed structure compared to interval ($t(38) = 2.485, p = 0.017$) as shown in Figure 2 C.

Discussion

We designed this study in order to investigate why and in which circumstances participants are pushed to drop a configural strategy, to explore and engage with alternative strategies. We designed a task where participants were not directly

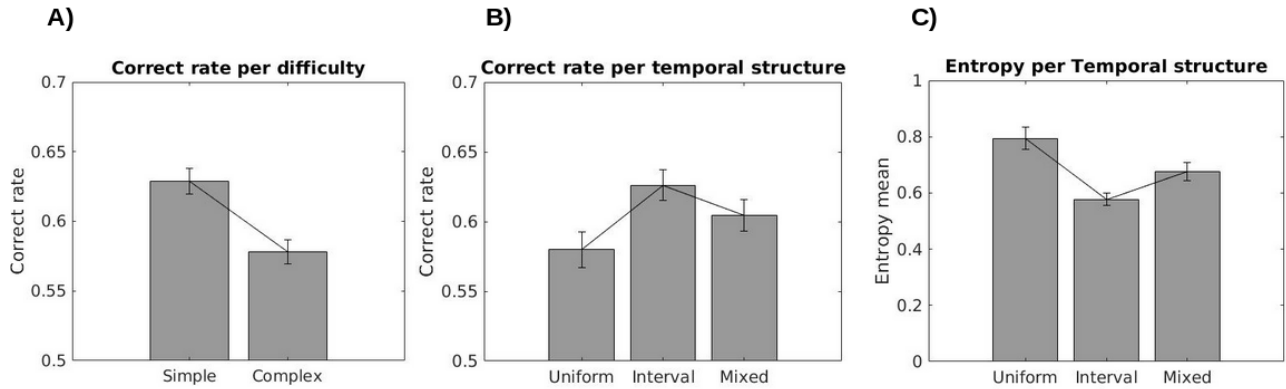


Figure 2: (A) Overall correct rate as a function of task difficulty. (B) Overall accuracy as a function of temporal structure. (C) The average entropy during simple and complex tasks as a functional temporal structure. Error bars are standard errors of the mean.

instructed to find a rule, but to learn the association between stimuli and outcomes. Participants had the possibility to explore different alternative strategies and use the one they preferred. Given the number of combinations among our stimuli (twenty-five), a configural strategy was cognitively expensive so the search for another strategy was implicitly incentivized. Strategy exploration comes with a cost though, so the trade off was between a cognitively expensive configural strategy or cognitively expensive strategy exploration.

We found that a high proportion of subjects correctly declared the underlying rule for the simple task (28/60) but not for the complex task (1/60), thus validating our task design. Participants who explicitly declared the strategy performed significantly better than those who did not.

As expected, a larger number of configurations than in Duncan et al. (twenty five versus six) pushed participants beyond the configural strategy. For the simple task the majority of subjects used the mixed configural/rule-based strategy. For the complex task, however, subjects continued to use the configural strategy. Strategy exploration was higher for the complex task, presumably because no explicit rule could be found.

We found main effects of temporal structure and task complexity on accuracy, but no interaction. This implies that the effect of temporal structure does not depend on difficulty. This may be a power issue, however, as tests for simple effects (interval versus uniform) were significant for the complex task but not for the simple task.

We had hypothesized that that temporal structure would have an effect on declarative knowledge. Specifically, that a mixed temporal structure (with focussed intervals followed by a revision period) would result in more subjects developing explicit knowledge. However this turned out not to be the case.

Our last set of findings was that strategy exploration was higher for the complex versus simple task and higher for the uniform temporal structure that is standard in WPT. This suggests that the more naturalistic temporal structures are less confusing for subjects.

Our results complement those found by Duncan et al. (2018). They showed how humans naturally learn relationships between outcomes and configurations of stimuli, and we have shown the contextual constraints on these behaviors. Our findings provide a foundation for further behavioural and neuroimaging studies. For example, will be able to use the Simple Task to study how explicit knowledge is derived from configural and mixed configural/rule-based strategies. Similarly, we can use the Complex Task to study purely implicit learning and use contrasts with the Simple Task to identify creation and use of explicit knowledge.

References

- Ballard, I., Miller, M., E., Piantadosi, T., S., & Goodman, D., N. (2017). Beyond reward prediction errors: Human striatum updates rule values during learning. *Cerebral Cortex*, *28*, 3965–3975.
- Domenech, P., & Koechlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current Opinion in Behavioral Sciences*, *1*, 101 - 106.
- Duncan, K., Doll, B., B., Daw, D., N., & Shohamy, D. (2018). More than the sum of its parts: A role for the hippocampus in configural reinforcement learning. *Neuron*, *98*, 645–657.
- Knowlton, J., B., Squire, R., L., & Gluck, A., M. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*, 106–120.
- Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A bayesian foundation for individual learning under uncertainty. *Front Hum Neurosci*, *5*, 39.
- Poldrack, A., R., Paré-Blagoev, A., R., Shohamy, D., Moyano, C., J., Myers, C., & Gluck, A., M. (2001). Interactive memory system in the human brain. *Nature*, *414*, 546–550.
- Speekenbrink, M., Channon, S., & Shanks, D. (2008). Learning strategies in amnesia. *Neuroscience and Biobehavioural Reviews*, *32*, 292-310.