

2019-09-09

Probing the early phase of rapid instructed rule encoding

Bugmann, G

<http://hdl.handle.net/10026.1/15143>

10.1016/j.biosystems.2019.103993

Biosystems

Elsevier BV

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Probing the early phase of rapid instructed rule encoding

Guido Bugmann^{a,*}, Jeremy Goslin^c, Serge Thill^{a,b}

^a Centre for Robotics and Neural Systems, Plymouth University, UK

^b Donders Institute for Brain, Cognition, and Behaviour, Radboud University Nijmegen, Netherlands

^c School of Psychology, University of Plymouth, UK

ARTICLE INFO

Keywords:

Instruction-based learning
RITL
Response time analysis
Error analysis
Neural mechanisms
Catch trials
Practice effects

ABSTRACT

Humans can rapidly convert instructions about a rule into functional neural structures used to apply the rule. The early stages of this encoding process are poorly understood.

We designed a stimulus–response (SR) task in which participants were first shown a SR rule on a screen for 200 ms, and then had to apply it to a test stimulus T, which either matched the S in the rule (SR trial) or not (catch trial). To investigate the early stages of rule encoding, the delay between the end of rule display and the onset of the test stimulus was manipulated and chosen between values of 50 ms to 1300 ms. Participants conducted three sessions of 288 trials each, separated by a median of 9 h. Random sequences of 20 rules were used. We then analysed the reaction times and the types of errors made by participants in the different conditions.

The analysis of practice effects in session 1 suggests that the neural networks that process SR and catch trials are at least partially distinct, and improve separately during the practice of respectively SR and catch trials. The rule-encoding process, however, is common to both tasks and improves with the number of trials, irrespective of the trial type.

Rule encoding shows interesting dynamic properties that last for 500 ms after the end of the stimulus presentation. The encoding process increases the response time in a non-stochastic way, simply adding a reaction time cost to all responses. The rule-retrieval system is functional before the encoding has stabilized, as early as 50 ms after the end of SR rule presentation, with low response errors. It is sensitive to masking however, producing errors with brief (100 ms) test stimulus presentations. Once encoding has stabilized, the sensitivity to masking disappears.

It is suggested that participants do encode rules as a parametrized function, using the same neural encoding structure for each trial, rather than reconfiguring their brain anew for each new SR rule. This structure would have been implemented from instructions received *prior* to the experiment, by using a library of neural functions available in the brain. The observed errors are consistent with this view.

1. Introduction

In recent years, the question of how task instructions are converted into neural representations that control behaviour has attracted a growing attention (Cole et al., 2013; Ruge and Wolfensteller, 2010, 2013; Ruge et al., 2018; Bugmann, 2009, 2012; Bugmann et al., 2013; Meiran et al., 2015; Palenciano et al., 2019). A better understanding of instruction-based learning (IBL) processes – also called “rapid instructed task learning” (RITL) – would have applications in education, medicine, social sciences and even robotics. A common observation in humans is that the internal encoding of a rule based on instructions is very fast. In particular, it is much faster than, e.g., learning from trials and errors.

How does the brain encode stimulus–response (SR) rules? In early work, we saw this as a theoretical problem of rapidly establishing a connection between sensory and motor parts of the brain, and considered a neural network model that rapidly connected a series of relays between sensory input and motor output (Bugmann, 2009, 2012). This was an early form of fast deep learning.

Further investigations into fast mechanisms of synaptic learning, combined with experimental data pointing to sub-second SR rule encoding times, suggested that the stimulus and the response neural representations should be separated by at most one synapse (Bugmann et al., 2013). Although we postulated such short links to be located in the hippocampus, data on the response time of human hippocampal neurons (Mormann et al., 2008) demonstrated that this route was too

slow to explain the experimental data.

I maging data also suggest that the medial temporal lobe (Hippocampus) is not involved in fast SR learning, but rather that the prefrontal cortex–premotor cortex (PFC–PM) complex is (Ruge and Wolfensteller, 2010; Cole et al., 2010, 2013). It is however surprisingly difficult to trace in the literature *direct* synaptic connections linking visually responsive prefrontal areas to motor areas. So, how can new rules be encoded in a short time when there is not enough time to set up multi-synaptic connections?

This puzzle led us to question whether it was necessary to re-wire a new network for each instruction. At least in experiments where large series of rules of a similar type are presented to participants, such as rules of the type “If the test stimulus is the letter A then press the right-side key, and if the test stimulus is not an A, press the SPACE bar”, a large number of combinations of letters and left/right responses can be generated for use in successive trials. Technically, participants will see a different rule in each trial and numerous instances of IBL will be measured. However, these rules are functionally very similar and differ only by the identity of the letter and the response. Therefore, the participant could also use a neural network architecture where only the identities of the test letter and response need encoding, e.g. memorizing the letter A and the specified response R. Such a rule-stimulus memory could then be compared with the actual test stimulus to decide whether to execute R or SPACE. The construction of such a network (IBL process) would typically happen during the instructions given to participants *prior* to the start of the experiment proper. In this case, we would not measure instances of IBL, but observe stimulus and response encoding effects.

The attractiveness of such a network is that the number of anatomical synaptic relays is not a limiting factor, because they do not need to be connected anew for each SR rule.

While it is difficult to test the fixed network hypothesis against, e.g. repeated re-construction, information can be gained from the dynamical properties of the biological network or the re-construction process in an experiment where processing time constraints are applied, e.g. by limiting the encoding time or limiting the duration of the test stimulus. One would expect a newly built network to not be able to function before all the connections are established, whereas a fixed network would only exhibit errors of stimulus and response encoding if tested at short notice.

Our new experimental design is of the type described above: $S \rightarrow R$ (SR trial) and not- $S \rightarrow \text{SPACE}$ (Catch trial). In half of the trials, the duration of the visual test stimulus is limited to 100 ms using a mask to probe the encoding progress at various times after the presentation of the rule. We initially used catch trials to maintain the focus of the participants; however their analysis turned out to provide additional information on the processes involved.

In the remainder of this paper, after presenting our methodology, we first give our results regarding reaction times, distributions of reaction times, and timing and type of erroneous responses. We also describe practice effects observed during the first session. In the discussion section, we then consider theories of task encoding such as look-up-table (LUT) learning of each SR instance, in the light of our results and show that these are made unlikely by the fact that catch trials add a huge number of rules to be learnt. Similarly, we find no strong support for the concept of brain network reconstruction for each rule. A more attractive possibility is that of a parametrized fixed network developed once at the start of the experiment. We propose a simple circuit diagram to capture its essence.

2. Methods

2.1. Experimental design

2.1.1. Overall design

The overall experiment consisted of repeated displays of SR rules followed by a test stimulus (see Fig. 1). First, the SR rule was presented

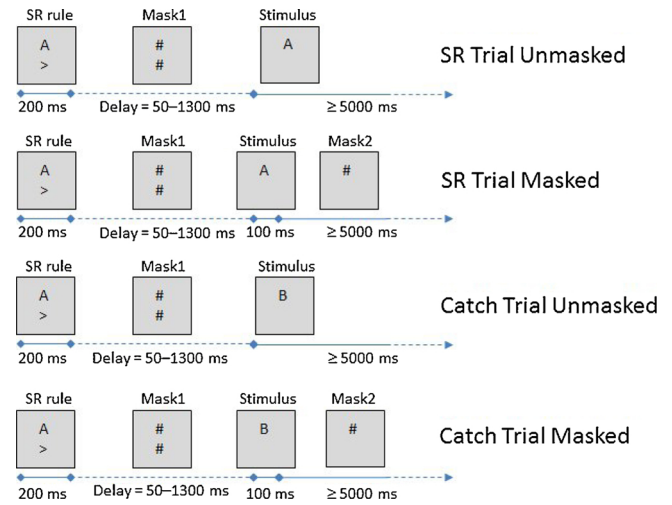


Fig. 1. The four experimental conditions used. The content of the greyed boxes is displayed on screen, within the foveal vision angle. (Note that the mask character (see text) is shown as “#” for convenience.)

visually for 200 ms. For example, the display “A>” was to be interpreted as “if you are shown an A, press the right-side key”. This was followed by a test stimulus T (e.g. “A”) after a delay. This delay was randomly set to one of eight possible values (50, 150, 300, 500, 700, 900, 1100 and 1300 ms) and the test stimulus persisted until a response was produced, or was masked after 100 ms.

Participants responded using “q” and “p” keys of the computer keyboard for “left” and “right” responses respectively (Fig. 2). The response was followed by a feedback screen displaying “correct” or “incorrect” and a progress bar to give a sense of progress through the experiment. The feedback display lasted until a key was pressed. Then a new trial started with a 1000 ms pause and a 1000 ms mask composed of two mask characters in the same positions as the symbols in the rule). The experiment was controlled by an eprime program.

2.1.2. SR and Catch trials, and masks

The test stimulus was the same symbol as in the SR rule (e.g., “A”) in two thirds of the trials (SR trials). In the remaining trials a different symbol (e.g., “B”) was displayed to that of the SR rule (Catch trials). Participants were required to press the space-bar when they encountered these Catch trials.




Fig. 2. Illustration of the position of hands and fingers during the experiment. Left and right index or middle fingers are positioned on the q and p key respectively. The two thumbs are placed on the space bar. Usually only one thumb is active during the experiment. The thumb is also used to trigger the start of the next trial. (Note that a full-size keyboard and large screen were used in the experiments.)

It is conceivable that participants use a serial strategy in performing the task, i.e. encode the rule first, then process the test stimulus. This approach is possible when the test stimulus is displayed until the response is produced. The only effect in trials with short delays would be to delay the response. So, in order to probe the state of the system *during* encoding, we limited the duration of the test stimulus and its persistence in the visual system by using a mask in 50% of the trials. In these *masked trials*, the test stimulus was shown for only 100 ms, after which it was replaced by the mask character (Fig. 1) (Rohenkohl et al., 2014).

2.1.3. Stimulus display

We used the 10 first capital letters of the alphabet as stimuli so that participants would not have to learn new stimuli as part of the task. The characters were of the type “Courier New”, size 40, bold, black on a white background. To complete the SR rule, the letter was randomly complemented with either a “>” or a “<”, as mentioned before (test stimuli consisted only of a letter). The mask was a composite character created by superposing all 10 characters used in the experiment. They subtended an angle of 1.34 degrees laterally, and the two characters one above the other in a rule subtended a vertical angle of 4.6 degrees. The display was a 22 in. LCD computer monitor with 1920 pixels by 1080 pixels.

2.1.4. Number of trials and sessions

For each of the eight possible delay values, there were 12 SR trials, 12 masked SR trials, 6 catch trials and 6 masked catch trials ($8 \times 36 = 288$ trials). As we used a pool of 20 SR rules (10 letters with two possible “<” and “>” responses to for each), participants were exposed to each rule on average $288/20 = 14.4$ times during a session. Participants thus saw both SR rules for a given letter, e.g. “A>” and “A<”, several times during a session.

In previous work, we noticed that the performance improved during the first session and then stabilized in subsequent sessions (Bugmann et al., 2013). We therefore ran three sessions in order to obtain stable performance data from sessions 2 and 3. Data from session 1 were used to investigate practice effects. The sessions lasted typically 30 min and were separated by time gaps ranging from 40 min to 7 days, with a median of 9 h.

2.2. Participants, instructions and practice trials

Participants were instructed that the two symbols displayed first were a S-R rule that had to be applied to the test stimulus displayed afterwards. They were told that if the test stimulus was different from the character in the rule, they should press the space-bar instead of producing the response specified in the rule. They were instructed to be as fast and accurate as possible. 30 practice trials were presented immediately prior to the first session of the experiment.

Participants were students and people from the general public. They were paid with course credits or shopping vouchers. 15 participants completed all three sessions.

Pre-processing of response times from sessions 2 and 3 led to the rejection of four participants, for the following reasons: (i) Participants were significantly slower than others in SR trials (e.g. average RT over all SR trials of 960 ms while all the others were less than 550 ms). (ii) Participants had several instances of large standard deviations (above 300 ms) in their response times, compared to an overall average SD of less than 100ms, (iii) Participants had large variations of RTs for neighbouring delay values (e.g. RT of 900 ms for delay = 900 ms and RTs around 600 ms for delays of 700 ms and 1100 ms). These criteria typically eliminated participants who were not engaging with the task or had frequent lapses of concentration.

The 11 retained participants were 6 females and 5 males, with an average age of 22.2 ± 7.3 . The number of participants was in line with previous experiments and provided sufficient statistics to distinguish the effects of different conditions explored here. The experimental design was approved by the Ethics Committee of the Faculty of Science of the University of Plymouth.

3. Results

In the following, we first examine practice trials, how RTs improved and errors diminished. The analysis of these changes as a function the overall number of trials and the number of trials devoted to SR trials and catch trials respectively reveals at least three distinct neural processes involved in the task. We then examined the differences in behaviour between short and long delays between SR rule presentation and rule retrieval. We find that the rule encoding process uniformly (not stochastically) adds time to all the responses. We then examine the effects of masking on response times and errors. We found that masking widens the distribution of SR response times, but not that of catch trials. In addition, the recognition of SR stimuli is intensity-sensitive, but not for catch trials. This provides further evidence for different processes being used these different trial types. Masking causes errors only for short delays after SR rule presentation, and these consist mainly of misclassifications.

3.1. Practice effects on response times

Fig. 3 shows (a) the average response time per delay is larger in session 1 compared to sessions 2 and 3. (b) That the response time is longer for short delays (50–300 ms) compared to long delays (700–1300 ms). We are going to examine the effect of practice on short-delay responses and long-delay responses separately. The assumption is that short delays probe the participant in a state of fresh or incomplete SR rule encoding. In contrast, for long delays, the rule encoding is complete and RTs reflect only on the retrieval processes.

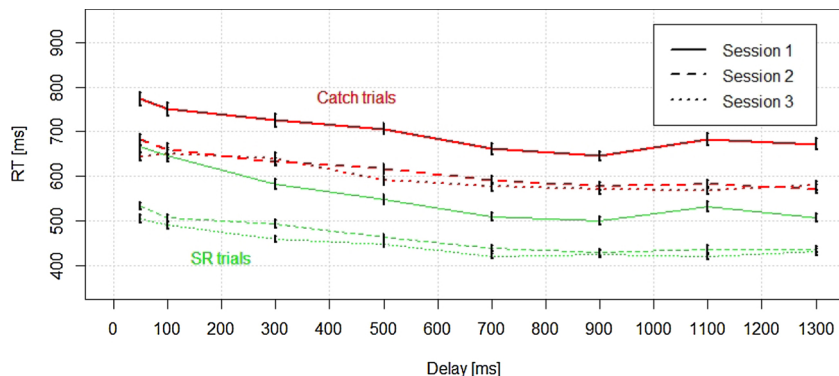


Fig. 3. Average RT for correct Catch and SR trials as a function of the delay between the end of the rule presentation and the start of the test stimulus, for all three sessions. Masked and non-masked trials are averaged together. Standard error bars are shown.




Fig. 4. Practice effects in session 1. The decrease of the average RT for four sets of data is shown. SR trials are shown in green (lower two curves), and catch trials in red (upper two curves). The dashed lines show short-delay data and the full line shows long-delay data. Masked and non-masked trials are averaged together. Standard error bars are shown. Note that for sessions 2 and 3, these four curves are horizontal lines.

To quantify practice effects, the RT data from the 288 trials in session 1 were grouped in 18 trial groups, each covering 16 trials (in other words, the groups 1–18 covered trials 1–16, 17–32, ..., 272–288 respectively).

Fig. 4 shows that long delays have RT diminishing almost linearly with the number of trials completed for both SR and catch trials. Short delays (dashed curves) have an RT dropping rapidly up until the halfway through the session (144 trials).

Long delay curves suggest an effect of trial type repetition, i.e. catch trials appear to improve more slowly, but a plot of the RT vs number of repetitions of trials of each type shows the same rate of improvement for catch and SR trials (see Fig. 5). This suggests a process-specific effect, i.e. catch processes and SR processes improve separately when they are practiced.

In contrast, the main effect on changes to short delay RTs was the number of trials practiced, with RTs stabilizing around halfway through the session. This suggests that the short-delay RTs are influenced by a process common to both types of trials, such as the encoding of the task.

If we assume that the RT at short delays include encoding, retrieval and motor costs, and the RTs at long delays include only retrieval and

motor costs, we can take the time difference between RTs for short and long delays in Fig. 4 as the encoding cost. We see that this stabilizes from about the middle of the session. Further RT reduction is then due to the processes active for long delays.

The encoding cost appears to be smaller for catch trials than for SR trials, but for sessions 2 and 3, i.e. in a practiced system, it is the opposite (50 ms for SR trials and 70 ms for catch trials (see Fig. 9). In a purely serial process, these two costs should be the same. Here, the retrieval process, during encoding, seems to be slightly longer for non-matching test stimuli. This may be caused by the additional classification process that is used in catch trials (see Section 3.4), but a detailed model will be needed to answer this.

3.2. Effect of practice on errors

We also examined the total number of errors over all participants as the sessions progress. Overall, our participants made very few errors, forcing us to use only 6 trial groups per session (thus covering 48 successive trials each). Since sessions two and three are qualitatively very similar, we only consider the average error rate for those (Fig. 6).

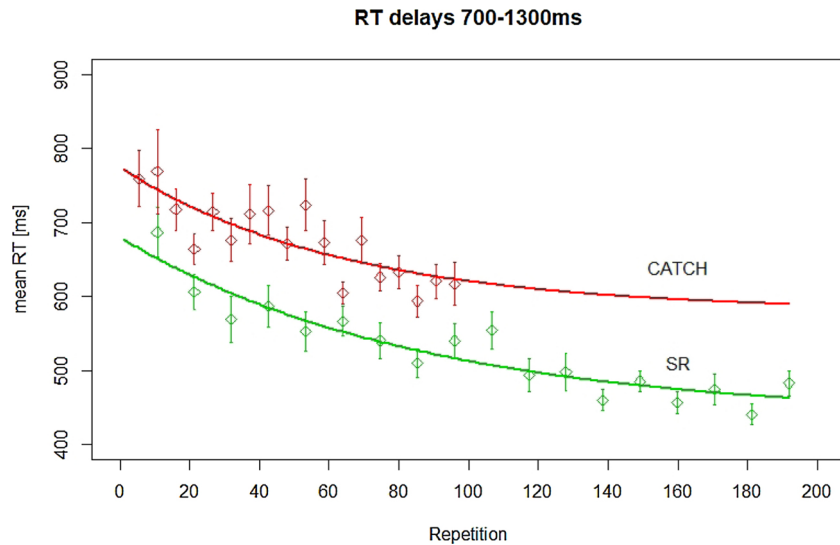


Fig. 5. Practice effects. Data for long delays in Fig. 4 are plotted as a function of the number of repetitions of each type of trial, SR and catch. Note that there are 20 individual SR rule and 180 individual catch rules (see discussion) that are not distinguished here.

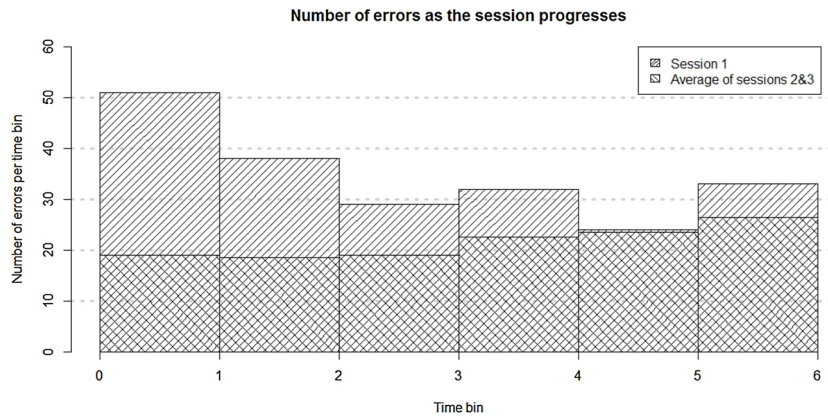


Fig. 6. Evolution of the number of errors as a session progresses. For sessions 2 and 3, the average number of errors is shown. There are no errors bars because, for session 1, the plot shows the counted number of errors, and for sessions 2 and 3, it shows the average between two counts.

We find that the number of errors diminishes during the first session and then stabilizes in session two and three, with maybe a hint of fatigue towards the end of sessions two and three.

In the first session, for all delays, two thirds of the errors occur during SR trials and one third during catch trials, reflecting the relative frequency of these trials. In sessions 2 and 3, however, half of the errors occur during SR trials, and half during catch trials (bearing in mind that the overall number decreased). This suggests a relative improvement of performance during SR trials than during catch trials, which is supportive of the idea that the two trials rely on distinct mechanisms.

What has changed? Lets us first define the type of errors that we observed: “SR_to_opposite” are errors where the participant responds with the opposite response, e.g. left instead of right, from the one defined in the SR rule. “SR_to_space” are errors where the participant presses the space-bar in a SR trial. “Catch_to_R” are errors where the participant produced the R response instead of pressing the space-bar. “Catch_to_opposite” are errors where the participant presses the response opposite to R instead of the space-bar.

If we now focus on the first 2 trial groups in session 1 and the first 2 trial groups in sessions 2 and 3 in Fig. 6 (this corresponds to the first 96 trials in these sessions), we find that SR_to_opposite errors have almost disappeared in sessions 2 and 3 (from 3.5% to 0.2% in masked and unmasked trials). SR_to_space errors dropped from 8.8% to 2% in masked trials and dropped slightly in unmasked trials (from 2% to 1.3%). Catch_to_R errors decreased barely in masked trials (6.3% to 5.6%) and increased in unmasked trials (from 1.7% to 3.7%).

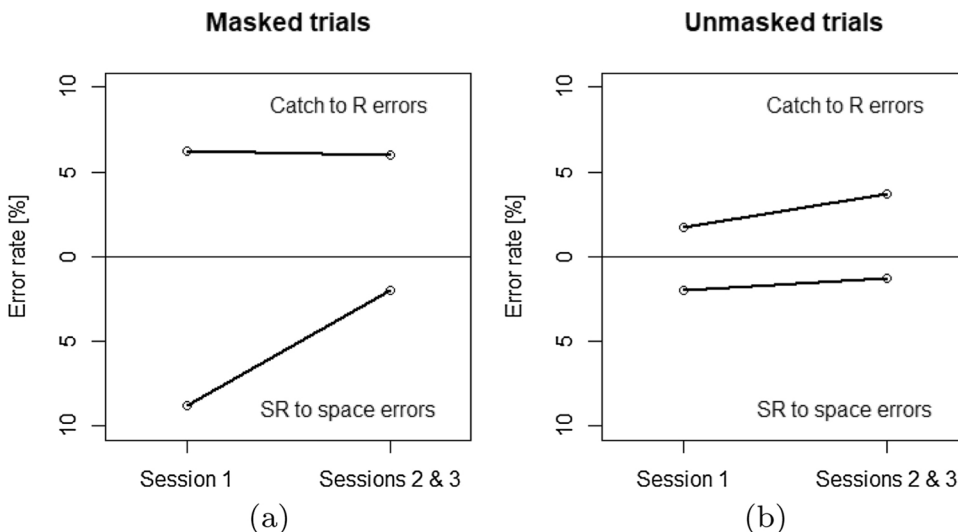


Fig. 7. Error rate during the first 96 trials of practice (session1) compared to the error rate during the first 96 trials of sessions 2 and 3. The bottom half of the figures has inverted Y-axes, so that points near the midline correspond to low error rates. The upper half represent errors where the response R from the SR rule is produced in a catch trial. The lower half represents errors where the space bar is pressed in a valid SR trial, i.e. where the test stimulus T is the same as the stimulus S specified in the SR rule.

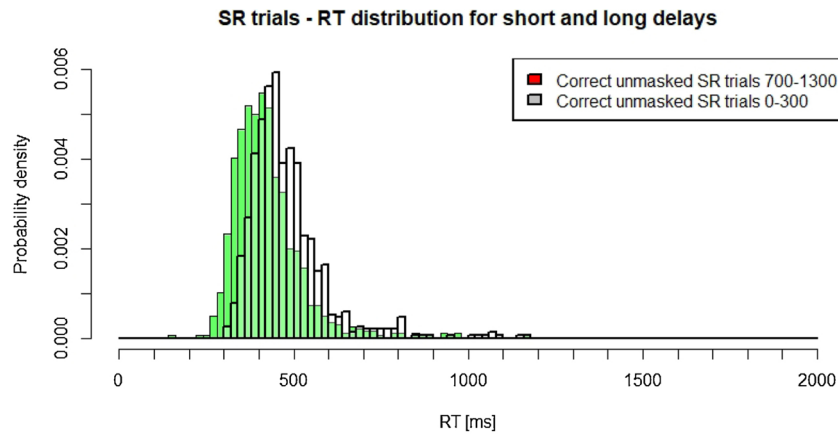
The results for R and space errors are summarized in Fig. 7. We ignored SR_to_opposite and Catch_to_opposite errors that became almost inexistent after practice. In these plots, a reduction of the distance between the two curves generally indicates an increase in selectivity, i.e. test stimuli T are more often correctly classified as “same as S” or “different from S”. This is visible in the masked trials, although there are also signs of an increased bias towards producing the R response. In the unmasked trials, there is no evidence for increased selectivity, but an increased tendency to produce R responses in error.

3.3. Effect of short delays on the distribution of responses times


We now focus on sessions 2 and 3 that show stable performances.

Fig. 3 (and Fig. 10a later on) show longer RTs for short delays between the end of the SR rule presentation and the start of the test stimulus. For sessions two and three, that RT cost is around 100 ms for the shortest delay and we see here that the encoding cost diminishes with the delay, fully disappearing for delays above around 500–700 ms. Here, we are interested in the mechanisms involved in the early phase of rule encoding (for delays from 50 ms to 300 ms).

First, we compare the distribution of RTs for unmasked correct SR trials for short delays (50–300 ms) vs long delays (700–1300 ms). Fig. 8 shows that the observed RT increase is due to a global shift of the distribution of RTs towards longer times. Q-Q plots (Fig. 9a) show that distributions for short and long delays have the same core width, and are shifted by around 50 ms to longer delays. There are also a small



(a)



(b)


Fig. 8. Short delay effects. Distribution of RTs for unmasked (a) correct SR trials and (b) correct Catch trials for delays 700–1300 (green or red) and 50–300 (translucent grey).

number of data points in a longer tail for short delays. This suggests that the process that is at work shortly after rule presentation is adding time, not variability. A few points of the tail of the long-delays distribution are moved even more, by around 120 ms. The behaviour is similar for catch trials, with a core shift of 70 ms, then a few long RTs shifted by 200 ms, from around 800 ms in long-delay RTs (Fig. 9b).

3.4. Effect of masking on RT distributions

Next, we note that masking increases slightly the average RTs for short delays (Fig. 10).

Looking at the effect of masking on the distribution of RTs for short delays in more detail, we find in (Figs. 11 and 12) different effects on SR



(a)

(b)

Fig. 9. Delay effects. Q-Q plots comparing the distributions of RTs for long and short delays for (a) SR and for (b) catch trials in sessions 2 and 3. The red line is the identity line. The blue line indicates by how much the distribution for short delays has been shifted to later times (by 50 ms for SR Trials, and 70 ms for catch trials) compared to the distribution for long delays. The light blue line indicates by how much a small number of points in the tail of the distribution for long delays have moved even further. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 10. Masking effects. Average RTs in sessions 2 and 3, for SR and catch trials, masked and unmasked. Standard error bars are shown.

and catch trials.

In SR trials, masking causes a widening of the distribution, revealed by larger-than-one slope in the $Q-Q$ plot (Fig. 12a). More precisely, masking does not change the distribution for unmasked RTs < 400 ms, presumably because these early RTs do not need T_s lasting longer than 100 ms. The longer unmasked RTs however are slowed down by the limitation in T duration. Then, a group of unmasked RTs between 700 ms and 800 ms correspond to masked RTs around 850–900 ms (flattening of the $Q-Q$ plot and a peak in the distribution of masked

trials), as if 850–900 ms was a preferred RT for a small set of longer responses. The cause is unclear. Are these “correct by chance” responses (see Section 3.6)? Note that masking does not have a statistical all-or-none behaviour here. Masked stimuli are processed, but seem to provide lesser “drive” for the triggering of the R responses.

In catch trials, the central parts of the distributions are the same (Fig. 12b). Some short RTs have disappeared due to masking and long ones have appeared. Here, masking does essentially not slow down responses (except for a group of late responses), indicating possibly that

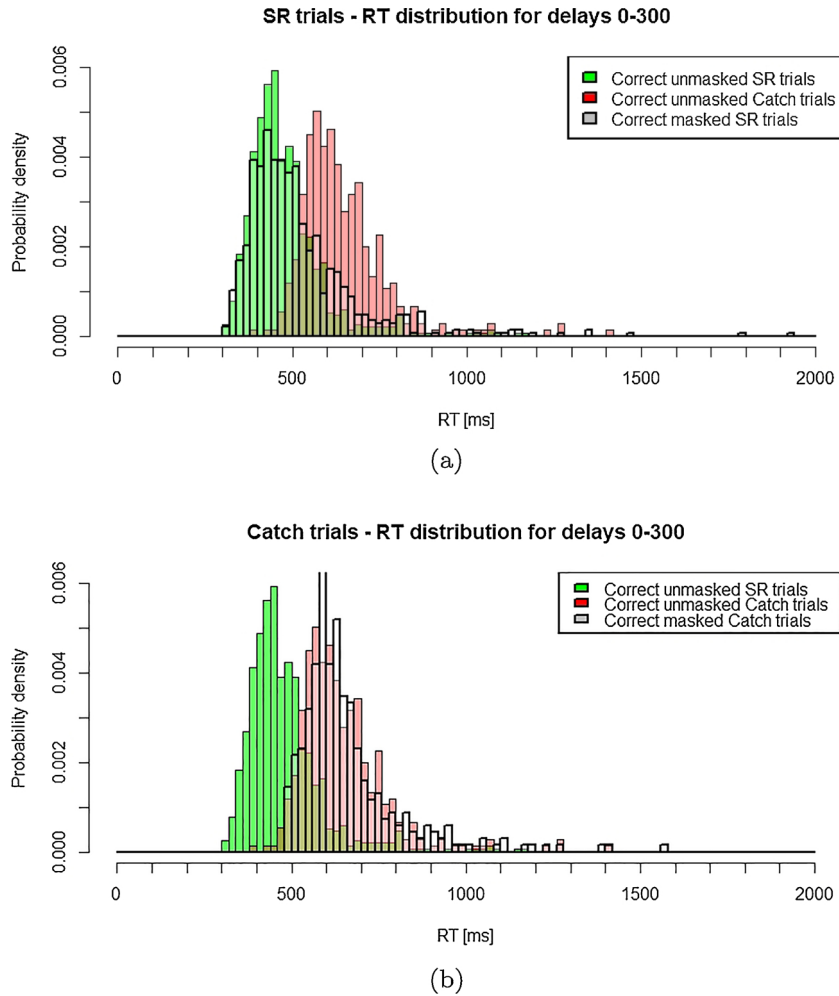


Fig. 11. Effects of masking. Histograms of response times for delays between 50–300 ms. The translucent bars represent (a) correct masked SR trials and (b) correct masked Catch trials. The green and red bars (left and right distributions) represent the correct responses in unmasked SR and catch trials respectively.

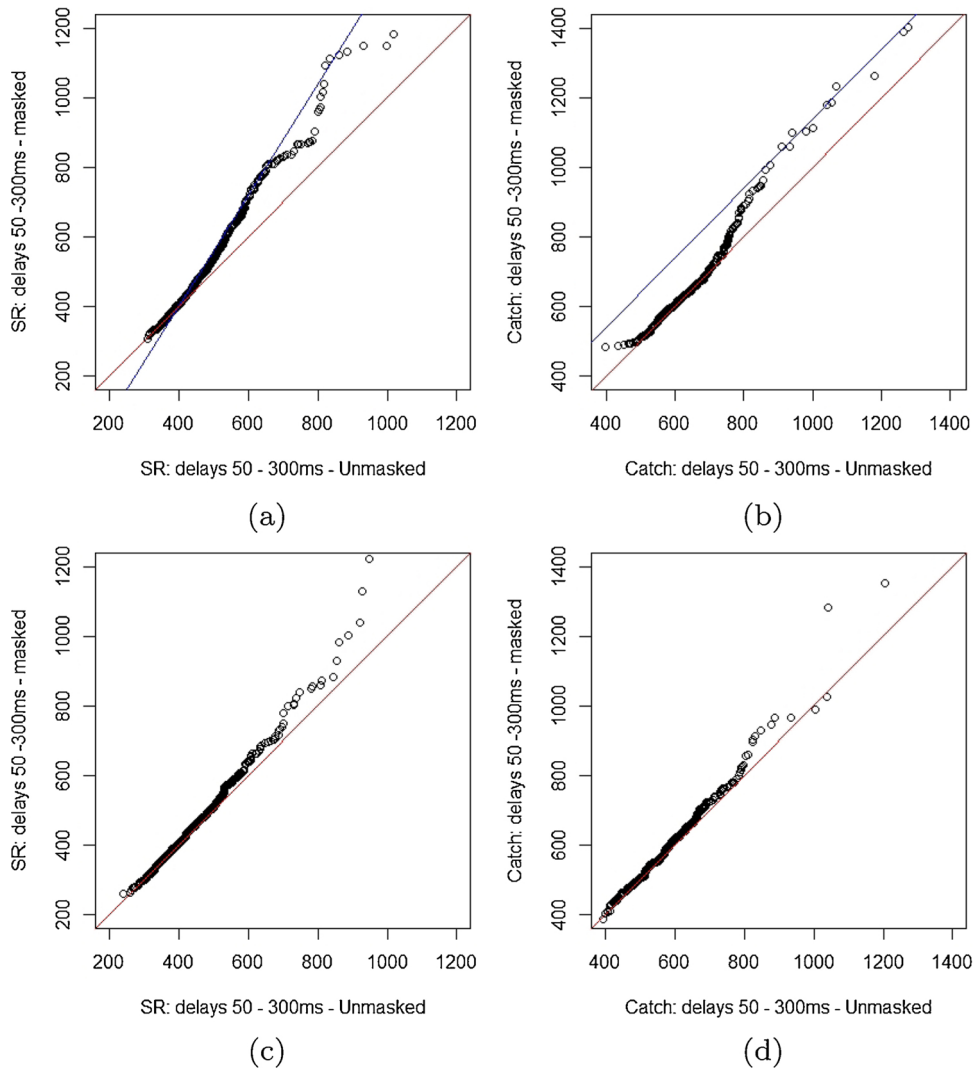


Fig. 12. Masking effects. Q-Q plots for short delays for masked vs unmasked trials for (a) SR trials and (b) catch trials. For comparison the long-delay data are also shown, for (c) SR trials and (d) catch trials. The red line is the identity. The blue line in (b) is shifted upwards by 140 ms. In (a) it has its origin at (400,400) and a slope of 1.6.

the decision about the mismatch between S and T takes an amount of time not influenced by the duration of T.

Thus, the generation of the R response seems directly driven by the T, while catch responses depend on a process for which a 100 ms stimulus is a sufficient input.

3.5. Effect of masking on errors

Masking increases the number of errors for short delays (Fig. 13). For unmasked trials, the number of errors stays low even for small delays, indicating that the rule encoding process has produced a




Fig. 13. Number of errors for each delay in masked and unmasked trials. (a) For session 1. (b) Average numbers over sessions 2 and 3. SR and catch trials show a similar behaviour and their counts are added up.

Table 1

Breakdown of responses into correct and errors (SPACE: space bar pressed; R: response specified in SR rule pressed; Opp R: opposite response pressed) for SR and Catch trials. Values highlighted with * are discussed in the text.

		Delay 50–300 ms		Delay 700–1300 ms	
		<i>n</i>	%	<i>n</i>	%
<i>SR trials</i>					
R Correct	No mask	768	97.34	1033	97.82
	Mask	740	94.39	1029	97.63
	Total	1508	95.87	2062	97.73
SPACE error	No mask	10	1.27	7	0.66
	Mask	30	3.83*	13	1.23
	Total	40	2.54	20	0.95
Opp R error	No mask	11	1.39	16	1.52
	Mask	14	1.79	12	1.14
	Total	25	1.59	28	1.33
<i>Catch trials</i>					
SPACE Correct	No mask	378	95.45	504	95.64
	Mask	346	87.37	504	95.45
	Total	724	91.41	1008	95.55
R error	No mask	17	4.29	20	3.80*
	Mask	48	12.12*	20	3.79*
	Total	65	8.21	40	3.79
Opp R error	No mask	1	0.25	3	0.57
	Mask	2	0.51	4	0.76
	Total	3	0.38	7	0.66

functional system already 250 ms after the start of the rule presentation.

Masking also affects the number and types of errors. To analyse these errors, we retain the division between a short (50–300 ms) and a long delay group (700–1300 ms). We compared the errors in masked and non-masked trials (see Table 1). For SR trials, we counted the number of SPACE responses produced in error, and the number of opposite responses (Opp R), e.g. “right” when a “left” was expected. For catch trials, we counted the number of responses (R) specified in the SR rule (thus produced in error), and the number of opposite responses (Opp R).

A number of observations can be made in Table 1. First, it is apparent that the percentage of correct results in unmasked SR trials is the same for short and long delays. There are more errors in catch trials, but here too there is no difference between short long delay, in unmasked trials. Thus, the process that increases the RTs for short delays does not affect the accuracy.

Catch trials tend to contain more errors than SR trials, and most of these errors are “R” type errors, *i.e.* producing the response indicated in the SR rule. This suggests that a priming mechanism exists for the response expected from the rule (Meiran et al., 2015). One could speculate that participants have treated the task as a “delayed- response” task with a R response prepared and using the test stimulus as a GO trigger (see, e.g. Wise et al., 1997).

The picture painted by the masked trials is, however, very different. Masking has no effect on errors for long delays (see e.g. the values highlighted with * in the bottom half of Table 1), but strongly increases the percentage of errors for short delays. Opposite R errors are almost inexistent in catch trials and seem triggered in SR trial with a probability independent on the delay or the mask. Since masking in itself does not affect the perception of the stimulus (because it has no such effect for long delays), it is not immediately clear how the mask might cause errors in short delay situations.

Note that a great majority of masked stimuli are processed correctly. In the masking literature, 100 ms is a relatively long stimulus duration that allows over 90% correct processing of the stimulus (see e.g. Bacon-Macé et al., 2005). According to Rohenkohl et al. (2014), masking

disrupts early iconic stimulus traces, with neurophysiological data collected from monkeys showing a weakening (less spikes per second) and shortening of the neural signal (Kovacs et al., 1995). Maybe this can cause a proportion of misclassification errors explaining the increases space errors in SR trials and increased R errors in catch trial.

3.6. Response time distribution of errors

To investigate how the mask might cause errors in more detail, we look at the RT distribution of error trials (Fig. 14). As before, we consider short and long delays separately, and are interested in differences between masked and unmasked trials for both catch and SR trials.

For unmasked trials (left column of Fig. 14), the errors appear all to be of the type “hasty misclassifications”, with early R produced in catch trials and early SPACE responses produced in SR trials (the number of errors taken from Table 1 is $N = 17$ for figure (a), $N = 10$ for (c), $N = 20$ for (e) and $N = 7$ for (g)). There is also a group of late RTs, above 1.1 s, that we tentatively identify as “I do not know” type errors.

In masked trials, only short delays show an effect of masking. In SR trials (Fig. 14d: $N = 30$) SPACE errors occur with RTs of normal SPACE responses, presumably caused by a T misclassification. In catch trials (Fig. 14b: $N = 48$), the picture is more complex. We find R produced with RTs early in the SR distribution, early in the catch distribution and some late in the catch distributions. These could correspond to an early bypass of the catch classification system, errors by the catch classification system and errors in response initiation, respectively. The mask thus seems to cause mainly erroneous classifications during encoding, and prevents hasty decisions.

4. Discussion

In this section we first discuss experimental evidences for distinct processes being involved in our task, and what data tell us about the dynamics of the rule encoding process. We then evaluate three theories of rule encoding: Look-up-tables of S-R rules, fast task switching between rule sets, and the use of a fixed parametrized network. The latter is suggested to be the more plausible for our experiment.

4.1. Practice effects, errors, delay effects and masking

Our task is not a straightforward SR rule encoding and retrieval task. The introduction of catch trials “to maintain the concentration of participants”, actually expands the rule-set significantly, from 20 to 200 (see Section 4.2.1 further down), and led us to analyse both SR and catch trials.

Analysing the practice session 1 reveals that processes involved in SR and catch tasks are, at least in part, distinct. The short-delays cost decreases with practice trials and reaches an asymptotic level for the same number of trials for SR and catch trials, suggesting that is not specific to either trial type. What comes to mind is that it reflects the development of the rule encoding process, given that there is no difference between SR and catch trials as far as rule encoding is concerned. It is only when the test stimulus is shown that the two trial types become different.

A slight problem is that the size of the encoding cost appears different for SR and catch trials (smaller for SR and larger for catch). This forces us to consider a more detailed model of encoding, bearing in mind that it may be the same neural system that is modified after rule presentation and that then classifies the test stimulus.

The analysis of the effects of masking on RT distributions suggests that SR and catch trials use a different decision process, at least during encoding, where the responses in SR trials is more directly driven by the test stimulus, whereas catch trials rely on an additional stimulus classification process that produces an output of constant magnitude. For long delays, both outputs appear to be independent on the stimulus

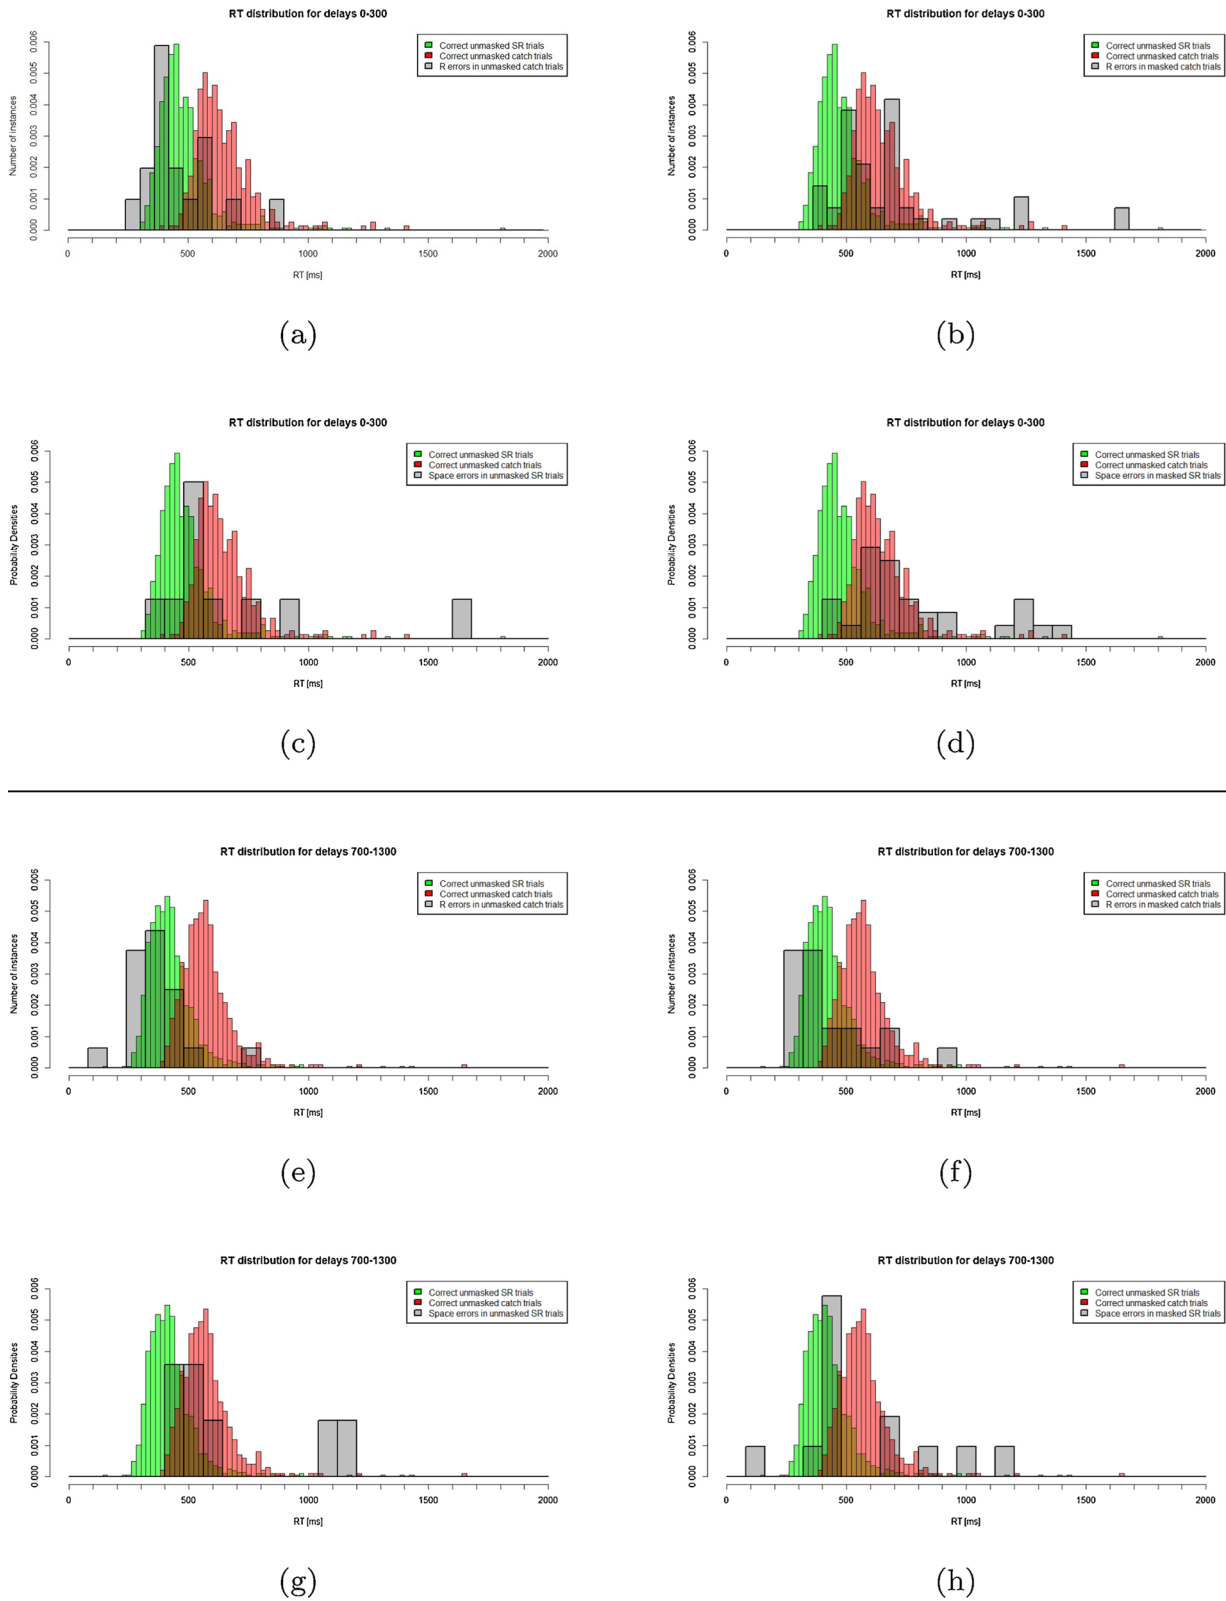


Fig. 14. Error distribution for short delays 50–300 ms (top half) and long delays 700–1300 ms (bottom half), for unmasked (left column) and masked trials (right column). The top row in each half (a, b; e, f) shows catch trial errors (responding R rather than pressing the space bar). The bottom row in each half (c, d; g, h) shows SR trial errors (space bar rather than responding R). The reference histograms in the background are for correct unmasked SR and catch trials for short delays (top half group) and long delays (bottom half group). Bins for errors are 80 ms except for 60 ms in (a, b) because the number of observations is small (see Table 1).

intensity.

For long delays, RTs decrease with trial type repetition, consistently with the idea that SR and catch response selection and execution are handled by distinct systems.

Our experiment used a time-constrained design to probe the early encoding stages. By examining sessions 2 and 3, that are free from practice effects, we found an encoding effect that slows down the responses by around 100 ms for both types of trials for the shortest delay (Fig. 3). This time is much shorter than the 500 ms delay range over which that effect is perceptible, suggesting a slow encoding process during which responses are delayed.

It is notable that the distribution of RTs for short delays is the same as for long delays, but shifted to longer RTs. Consistently, the standard deviation of RTs is not increased for short delays (see e.g. Fig. 3). Thus, the encoding process delays all responses by an amount of time that is largest in its early stages. One can think of a dynamical system that is initially imprinted with the SR rule and takes some time (e.g. 500 ms) to become fully stabilized, during which it has less power to trigger a response to a test stimulus. The diffusion framework nicely explains response delays when the input signal is weak (see e.g. Ratcliff et al., 2016).

Masked trials show that the labile early encoding state finds it difficult to process brief visual information and makes a number of mainly classification errors (for long delays, masking causes no problems at all, and a 100 ms stimulus is perfectly sufficient for rule retrieval). Instead, during early encoding, a more persistent visual input is needed to generate a decision of match or non-match. This is another feature that a detailed model should reproduce.

4.2. Task encoding

We examine here different methods by which the rule-encoding process could take place: exemplar learning, network re-configuration, or using a fixed network with dynamic elements.

4.2.1. Exemplar learning

This section covers a process called sometime exemplar learning, associative retrieval or look up table (LUT) approach. Indeed, participants could avoid encoding rules by simply memorizing all instances as input–output mappings. For instance one of our 20 SR rules, e.g. the “A >” rule, could be converted into 10 associations (we use a set of 10 letters) where the first represents the SR rule and the 9 others represent the catch rule:

“A >” followed by A → respond right;

“A >” followed by B → respond SPACE;

...

“A >” followed by J → respond SPACE;

Thus, our experiment could be represented by $20 \times 10 = 200$ input–output mappings (20 SR mappings and 180 catch mappings). In Forrester et al. (2014), associative input–output learning took around 75 trials per mapping (their experiment used a task cue followed by one of four numbers to classify either as low/high or even/odd, as specified by the task cue. There were 8 possible combinations). In Palmeri (1997), around 200 trials are needed per mapping to convert a counting process into an input–output mapping (participants see a card with N randomly placed dots and need to press the Nth key. There were 5 examples for 6 possible numbers, i.e. 30 combinations).

In our experiment most of the LUT learning would have happened in session 1, as we see no practice effects in sessions 2 and 3. The 20 SR rules are used in 192 trials, i.e. each SR mapping is seen 9.6 times in session 1 (a number similar to the rule memorizing experiment by Ruge and Wolfensteller (2010) – using only 4 rules). However, there are $20 \times 9 = 180$ catch rules but only 96 catch trials. Hence each catch input–output mapping is only seen in average 0.533 times in session 1.

With such small number of repetitions, exemplar learning does not seem practically possible here, at least for the catch task. For such cases,

learning the task rather than exemplars appears to be a more feasible solution. Different considerations led Forrester et al. (2014) to a similar proposal: “...associatively-based compound retrieval can serve as the basis for performance with a small stimulus set. But, when organization by tasks is apparent, control via task-set selection is the natural and efficient strategy”.

4.2.2. Network re-configuration

What we normally understand by Rapid Instructed Learning, is the rapid generation of a “neural program” that can execute the instructed task. However, little is known about the mechanisms of the corresponding task encoding. In some cases, participants report silent rehearsal, but it is unknown how this leads to readiness to respond when a test stimulus is presented. Ruge and Wolfensteller (2010) propose a process by which “abstract” rules are converted into “pragmatic” (executable) ones. Cole et al. (2017) suggest that the frontoparietal network consists of flexible hubs that help implement flexible changes in cognitive programs during RITL. It is proposed that such a program generates responses as prepared reflexes (Meiran et al., 2015). A number of papers deal with “cognitive control” that “...proactively configures information processing to suit expected task demands” (Jiang et al., 2018). Typical experiments to examine cognitive control use task switching protocols. These use a number of tasks that participants know in advance and give them cues as to what task to perform in a given trial. One could argue that such experiments do not test the conversion of instructions into task representations, but only the recall of a task-set.

Task switching causes response time costs of over 150 ms at short retrieval delays (Monsel, 2003; Hunt and Klein, 2002). That cost diminishes with practice, leading Meiran (1996) to suggest that practice led to faster reconfiguration. Experiments by Strobach et al. (2013) and Pashler and Baylis (1991) show that practice speeds up the “central response-selection” stage (not perceptual or motor stages). Hunt and Klein (2002) note that switching costs disappear for long delays.

There are many similarities between these observations and our results, such as an RT cost for short delays, with a reduction of this cost with practice, and elimination of any cost for long delays. The only difference is that task switching increases errors at short delays (Monsel, 2003), while our error rate remains constant for all delays. It is unclear whether practice effects are similar, as we have not found an analysis similar to ours for task switching experiments.

Is our experiment a “task switching” experiment? In Monsel (2003), tasks are functions applied to the test stimulus T. e.g. function F1(T) tells if a test number odd or even, function F2(T) tells if it is large or small. Task switching is function switching. In our experiment we have a set {A, B, ..., J} of 10 possible values for T, and in each trial we use a different function F(T) corresponding to a different SR rule, e.g. if T = A, then press p, or if T = B press SPACE, or if T = C, press SPACE, ... So, in that sense there is task switching in each trial, but it requires each of the 20 tasks to be encoded in advance. A difficulty is that we did not instruct participants of each individual task and there is not enough exposure during practice to learn them.

Another view would be that participants use only one function F(SR, T), where the SR rule is a parameter of the function, and it is acquired and stored in each trial. This view would not apply to the task in Monsel (2003), because the odd/even and large/small functions do not just differ by some parameters.

Therefore we are tempted to paraphrase Forrester et al. (2014) by suggesting that when task switching involves parametrized functions, the natural and efficient strategy is to implement the function once at the start of the experiment, then modify only the parameters for each new task.

In support for this idea is the lack of dysfunctional period after the presentation of a new rule. We speculated in the introduction that participants would not be able to execute the task before re-configuration has taken place, but they were able to generate error-less

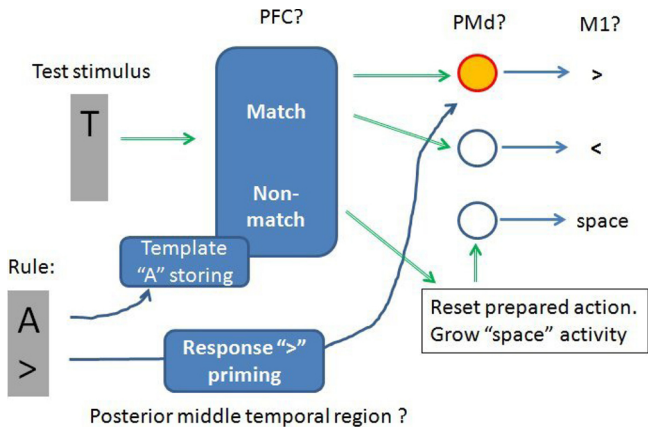


Fig. 15. Hypothetical network linking a set functions from repertoires in the PFC and the PM to perform our SR learning task. The anatomical labels are speculative (see text). PFC: premotor cortex; PMd: premotor area dorsa; M1: primary motor cortex.

responses to test stimuli presented 250 ms after the start of the display of the new rule. Yes, masking shows an early-stage difficulty in classifying the test stimulus but the very few resulting erroneous responses could only be produced by a functional network. Indeed this argumentation would have to be revised if new data show that a functional network can be built anew in less than 250 ms.

4.2.3. Fixed network

Hommel (2000) noted that the structure of instructions influence the organization of the task. Lets assume that participant literally encoded the instruction given prior to the measurements: “ If the $T = A$, respond with the R in the rule. If it T different from A, press the space bar”.

The brain provides a library of functions, or “repertoire”, that can help design such a fixed network. For instance, neurons in prefrontal cortex (PFC) are known to encode whether two successive stimuli match or do not match. Note that there is an ongoing debate whether this function is permanently available (Hussar and Pasternak, 2013; Qi et al., 2012), or is built in a task-specific way (Stokes et al., 2017). Similarly, the motor system has a number of neurons that trigger specific actions, such as the finger presses needed in this experiment (Graziano and Aflalo, 2007).

Using such repertoire-functions, the task could be implemented as a fixed network, e.g. as in Fig. 15, in which SR-rule encoding can be achieved with no synaptic learning at all. Such a network would also provide an efficient solution to the encoding of catch trial responses, by associating the “not match” function to the space-bar response, instead of encoding 9×10 catch rules.

Two parameters would need memorizing: (i) the nature of the response R and (ii) the identity of the stimulus S to compare to the future test stimulus T.




Fig. 16. Estimated error rates in SR and catch trials, in (a) the unmasked case, and (b) the masked case. The match detection box represents the process of deciding whether the test stimulus is the same as in the SR rule. The Non-Match detection box does the opposite. The fast process triggers the primed response R. The slow process needs to re-program the response from R to SPACE.

These memorizing processes should only affect the responses for short delays. Response storage can be very fast, by starting sustained firing in some neurons. There are data on early response encoding, e.g. in the premotor cortex (PM) (see e.g. Ledberg et al., 2007). The short delay effects would therefore be linked mainly to the S memorizing processes. We mentioned earlier a dynamic systems hypothesis of a combined memorizing-and-comparing mechanism. A more detailed model of that process would be needed to explain the short-delay effects.

With a fixed network, practice would improve its function, instead of the individual SR mappings. One can speculate that the sub-functions of the network are only improved when they are practiced. The memorizing stage is common to all trial types and improves with the number of trials performed, as revealed by a reduction of the short-delay cost (Section 3.1). The functions processing match and non-match events would improve separately, when respectively SR and catch trials are practiced, as revealed by a reduction of the RTs for long-delay trials during practice. As a side note, it is striking that practice effects completely disappear in sessions 2 and 3 (graphs not shown).

In the fixed network model, the RT difference between SR and catch trials (i.e. with an average of 153 ms, and no systematic variation with the delay) is likely to be due to the differences in processing speed in the two distinct processing streams. One hypothesis is that the R response is primed during SR presentation, hence is executed as a “prepared reflex” (Meiran et al., 2015; Muhle-Karbe et al., 2016). Whereas, in the case of a catch trial, the prepared response needs first to be inhibited and then the SPACE response ramped up (Pouget et al., 2017). There is also a possible extra classification process in catch trials (Section 3.4).

4.3. Explaining errors using the fixed network model

Returning to the types of errors made by participants, we can note that they do easily map onto various stages of the network (Fig. 16) based on the details of their RT distributions discussed before. First a test stimulus can be misclassified and the wrong stream activated. This generates mainly R or SPACE responses with the normal RT of these responses.

Another type of error occurs at the stage of action initiation, where the re-programming from R to SPACE may fail and the R response is produced in a catch trial. Based on the timing and count of errors in Fig. 13 and Table 1, we have estimated rates for this error for short delays (see Fig. 16).

The estimates in Fig. 16 indicate that one can expect masking-induced miss-classification in favour of a “match”, and a weakening of response re-programming in catch trials.

5. Conclusion

We have described an experiment of fast SR rule encoding and probed the time course of the process by presenting test stimuli at

various delays after the rule specification.

The system stabilizes 700 ms after the start of the SR rule presentation, but is functional 250 ms after the start of presentation. The data suggest that, while rule encoding progresses, the system is already able to process test stimuli, albeit slower in the early stages, and with some weakness when faced with brief visual stimuli. It will be interesting and challenging to design a more detailed dynamical model that replicates these properties.

There are puzzling facts for the model to explain: During encoding, T processing in SR trials is slowed down by masked stimuli, but not in catch trials. However, the encoding cost in unmasked catch trials is larger than in SR trials.

The “simple” SR rules used in the experiment actually require a complicated task encoding, due to the presence of catch trials. It appears that the simplest way for participants to perform the experiment is to encode a “parametrized function” at the start of the experiment, then use the SR rules as parameters, rather than attempting to re-encode the rules anew in each trial. There is a library of functions in the brain that are convenient to build a network implementing such a parametrized function.

Such a “fixed” network is consistent with stream-specific practice effects observed in session 1, the characteristics of the observed errors and the analysis of RT distributions.

Conflict of interest

The authors declare no conflict of interest.

Acknowledgments

We are grateful to Frank Loesche for creating the composite mask character and running several of the experiments.

References

- Bacon-Macé, N., Macé, M.J.-M., Fabre-Thorpe, M., Thorpe, S.J., 2005. The time course of visual processing: backward masking and natural scene categorisation. *Vis. Res.* 45 (11), 1459–1469.
- Bugmann, G., 2009. A neural architecture for fast learning of stimulus–response associations. In: Prasad, B., Lingras, P., Ram, A. (Eds.), *Proceedings of the 4th Indian International Conference on Artificial Intelligence, IICAI 2009*. Tumkur, Karnataka, India, December 16–18. pp. 828–841.
- Bugmann, G., 2012. Modeling fast stimulus–response association learning along the occipito- parieto-frontal pathway following rule instructions. *Brain Res.* 1434, 73–89.
- Bugmann, G., Goslin, J., Duchamp-Viret, P., 2013. The speed of learning instructed stimulus–response association rules in human: experimental data and model. *Brain Res.* 1536, 2–15.
- Cole, M.W., Bagic, A., Kass, R., Schneider, W., 2010. Prefrontal dynamics underlying rapid instructed task learning reverse with practice. *J. Neurosci.* 30 (42), 14245–14254.
- Cole, M.W., Braver, T.S., Meiran, N., 2017. The task novelty paradox: flexible control of inflexible neural pathways during rapid instructed task learning. *Neurosci. Biobehav. Rev.* 81, 4–15 The power of instructions: the influence of instructions on cognition, behaviour and physical states.
- Cole, M.W., Laurent, P., Stocco, A., 2013. Rapid instructed task learning: a new window into the human brain's unique capacity for flexible cognitive control. *Cogn. Affect. Behav. Neurosci.* 13 (1), 1–22.
- Forrest, C.L., Monsell, S., McLaren, I.P., 2014. Is performance in task-cuing experiments mediated by task-set selection or associative compound retrieval? *J. Exp. Psychol.: Learn. Mem. Cogn.* 40 (4), 1002–1024.
- Graziano, M.S., Aflalo, T.N., 2007. Mapping behavioral repertoire onto the cortex. *Neuron* 56 (2), 239–251.
- Hommel, B., 2000. The prepared reflex: automaticity and control in stimulus–response translation. *Control of Cognitive Processes: Attention and Performance Vol. XVIII*. MIT Press, Cambridge, MA, pp. 247–273.
- Hussar, C.R., Pasternak, T., 2013. Common rules guide comparisons of speed and direction of motion in the dorsolateral prefrontal cortex. *J. Neurosci.* 33 (3), 972–986.
- Hunt, A.R., Klein, R.M., 2002. Eliminating the cost of task set reconfiguration. *Mem. Cogn.* 30 (4), 529–539.
- Jiang, J., Wagner, A.D., Egner, T., 2018. External and internally generated task predictions are integrated in prefrontal cortex to jointly guide cognitive control. *Elife* 7, e39497.
- Kovacs, G., Vogels, R., Orban, G.A., 1995. Cortical correlate of pattern backward masking. *Proc. Natl. Acad. Sci. U. S. A.* 92 (12), 5587–5591.
- Ledberg, A., Bressler, S.L., Ding, M., Coppola, R., Nakamura, R., 2007. Large-scale visuomotor integration in the cerebral cortex. *Cereb. Cortex* 17 (1), 44–62.
- Meiran, N., 1996. Reconfiguration of processing mode prior to task performance. *J. Exp. Psychol.: Learn. Mem. Cogn.* 22 (6), 1423–1442.
- Meiran, N., Pereg, M., Kessler, Y., Cole, M.W., Braver, T.S., 2015. The power of instructions: proactive configuration of stimulus–response translation. *J. Exp. Psychol.: Learn. Mem. Cogn.* 41 (3), 768–786.
- Monsell, S., 2003. Task switching. *Trends Cogn. Sci.* 7, 134–140.
- Mormann, F., Kornblith, S., Quiroga, R.Q., Kraskov, A., Cerf, M., Fried, I., Koch, C., 2008. Latency and selectivity of single neurons indicate hierarchical processing in the human medial temporal lobe. *J. Neurosci.* 28 (36), 8865–8872.
- Muhle-Karbe, P.S., Duncan, J., Baene, W.D., Mitchell, D.J., Brass, M., 2016. Neural coding for instruction-based task sets in human frontoparietal and visual cortex. *Cereb. Cortex* 27 (3), 1891–1905.
- Palenciano, A.F., González-García, C., Arco, J.E., Ruz, M., 2019. Transient and sustained control mechanisms supporting novel instructed behavior. *Cereb. Cortex* 29 (9), 3948–3960. <https://doi.org/10.1093/cercor/bhy273>.
- Palmeri, T.J., 1997. Exemplar similarity and the development of automaticity. *J. Exp. Psychol.: Learn. Mem. Cogn.* 23 (2), 324–354.
- Pashler, H., Baylis, G.C., 1991. Procedural learning: I. Locus of practice effects in speeded choice tasks. *J. Exp. Psychol.: Learn. Mem. Cogn.* 17, 20–30.
- Pouget, P., Murthy, A., Stuphorn, V., 2017. Cortical control and performance monitoring of interrupting and redirecting movements. *Philos. Trans. R. Soc. B: Biol. Sci.* 372 (1718), 20160201.
- Qi, X.-L., Meyer, T., Stanford, T.R., Constantinidis, C., 2012. Neural correlates of a decision variable before learning to perform a match/non-match task. *J. Neurosci.* 32 (18), 6161–6169.
- Ratcliff, R., Smith, P.L., Brown, S.D., McKoon, G., 2016. Diffusion decision model: current issues and history. *Trends Cogn. Sci.* 4, 260–281.
- Rohenkohl, G., Gould, I.C., Pessoa, J., Nobre, A.C., 2014. Combining spatial and temporal expectations to improve visual perception. *J. Vis.* 14 (4), 8.
- Ruge, H., Karcz, T., Mark, T., Martin, V., Zwosta, K., Wolfensteller, U., 2018. On the efficiency of instruction-based rule encoding. *Acta Psychol.* 184, 4–19.
- Ruge, H., Wolfensteller, U., 2010. Rapid formation of pragmatic rule representations in the human brain during instruction-based learning. *Cereb. Cortex* 20 (7), 1656–1667.
- Ruge, H., Wolfensteller, U., 2013. Functional integration processes underlying the instruction-based learning of novel goal-directed behaviors. *Neuroimage* 68, 162–172.
- Stokes, M.G., Buschman, T.J., Miller, E.K., 2017. Dynamic coding for flexible cognitive control. *The Wiley Handbook of Cognitive Control*. John Wiley & Sons, Ltd., pp. 221–241 (Chapter 13).
- Strobach, T., Liepelt, R., Pashler, H., Frensch, P.A., Schubert, T., 2013. The time course of visual processing: backward masking and natural scene categorisation. *Atten. Percept. Psychophys.* 75, 900–920.
- Wise, S.P., Boussaoud, D., Johnson, P.B., Caminiti, R., 1997. Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. *Annu. Rev. Neurosci.* 20 (1), 25–42.