



City Research Online

City, University of London Institutional Repository

Citation: Cooper, R.P., Cook, R., Dickinson, A. & Heyes, C.M. (2012). Associative (not Hebbian) learning and the mirror neuron system. *NEUROSCIENCE LETTERS*, 540, pp. 28-36. doi: 10.1016/j.neulet.2012.10.002

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/4540/>

Link to published version: <https://doi.org/10.1016/j.neulet.2012.10.002>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Submitted to *Neuroscience Letters*, special issue ‘The Mirror Neuron System: A Bottom-Up Approach’, edited by Antonino Casile

ASSOCIATIVE (NOT HEBBIAN) LEARNING
AND THE MIRROR NEURON SYSTEM

Cooper¹, R. P., Cook², R., Dickinson³, A., & Heyes⁴, C. M.

¹Department of Psychological Sciences, Birkbeck, University of London

²Department of Psychology, City University, London

³Department of Experimental Psychology, University of Cambridge

⁴All Souls College & Department of Experimental Psychology, University of Oxford

Corresponding Author:

Richard P. Cooper

Department of Psychological Sciences,

Birkbeck, University of London

Malet Street, London WC1E 7HX

EMAIL: R.Cooper@bbk.ac.uk

PHONE: +44 20 7631 6211

Abstract

The associative sequence learning (ASL) hypothesis suggests that sensorimotor experience plays an inductive role in the development of the mirror neuron system, and that it can play this crucial role because its effects are mediated by learning that is sensitive to both contingency and contiguity. The Hebbian hypothesis proposes that sensorimotor experience plays a facilitative role, and that its effects are mediated by learning that is sensitive only to contiguity. We tested the associative and Hebbian accounts by computational modelling of automatic imitation data indicating that MNS responsivity is reduced more by contingent and signalled than by non-contingent sensorimotor training (Cook, Press, Dickinson, & Heyes, 2010). Supporting the associative account, we found that the reduction in automatic imitation could be reproduced by an existing interactive activation model of imitative compatibility when augmented with Rescorla-Wagner learning, but not with Hebbian or quasi-Hebbian learning. The work argues for an associative, but against a Hebbian, account of the effect of sensorimotor training on automatic imitation. We argue, by extension, that associative learning is potentially sufficient for MNS development.

Highlights

- Mirror neuron system responsivity is changed by sensorimotor experience
- Contingent sensorimotor experience is more effective than non-contingent or signalled
- Computational modelling indicates this is due to associative, not Hebbian, learning
- Associative learning (Rescorla-Wagner model) depends on prediction error
- Associative, but not Hebbian, learning is potentially sufficient for MNS development

Keywords: automatic imitation; associative learning; Hebbian learning; mirror neuron system; interactive activation; computational model.

Introduction

There is now ample evidence that experience is important in the ontogeny of the mirror neuron system (MNS). We know that monkeys can develop ‘tool-use mirror neurons’ [11]; expertise in an action domain selectively enhances the responsivity of the human MNS to actions in that domain (music – [14]; dance – [1]); and laboratory-based training can both increase [23] and decrease [4, 12] MNS responsivity. However, three related questions about the ontogeny of the MNS remain to be resolved: 1) Does experience play a facilitative ‘tuning’ role or an inductive ‘forging’ role? It would be tuning if experience modulates the rate or specificity with which the MNS acquires the capacity to map observed onto executed actions, but the MNS would eventually develop visual-motor matching properties even in the absence of experience. It would be forging if experience is *necessary* for the development of MNS—that is, if, in the absence of experience, neurons in the inferior parietal and premotor cortex areas would not become responsive to the sight, as well as the performance, of certain actions. 2) What type of experience is crucial? Does the development of the MNS depend on seeing actions (sensory experience), on performing actions (motor experience), and/or on correlated observation and execution of the same actions (sensorimotor experience)? 3) Does sensorimotor experience contribute to the development of the MNS via associative or Hebbian learning? Associative learning depends on contingency as well as contiguity; the connection between two neurons or event representations is strengthened if they are activated at about the same time *and* activation of one is relatively the best predictor of activation of the other. In contrast, Hebbian learning depends on contiguity alone. This article is concerned primarily with the third and most specific of these questions, but as we shall see, the answer to this question has implications with respect to the other, broader questions about MNS development.

Catmur [2] has reviewed the evidence that sensorimotor experience makes a distinctive contribution to the development of the MNS. Behavioural, electrophysiological and neuroimaging studies with adult humans have shown that, when the potential contributions of sensory experience and motor experience are controlled, a relatively brief period of novel sensorimotor experience can change MNS responsivity in a variety of ways. ‘Compatible’ sensorimotor training, where action execution is paired with the observation of similar actions, can enhance mirror responses; ‘incompatible’ sensorimotor experience, where action execution is paired with observation of dissimilar actions, can reduce, abolish or even reverse mirror responses; and ‘arbitrary’ sensorimotor training, where action execution is paired consistently with simple colours or shapes, can induce the MNS to respond to inanimate stimuli.

Many of the experiments indicating that sensorimotor experience can change the functioning of the MNS were designed to test the associative sequence learning model (ASL [16, 17]). This model proposes that the development of the MNS is mediated by the same, phylogenetically ancient mechanisms of associative learning that produce Pavlovian and instrumental conditioning. Studies of human and nonhuman animals have shown that these mechanisms of associative learning are sensitive, not only to the contiguity between events (i.e. how closely together they occur in time), but also to the contingency, or predictive relationship, between events. Therefore, associative learning not only increases with the probability of the second event (E_2) given the first event (E_1), or in other words with the likelihood of contiguous pairings of E_1 and E_2 , but also decreases with the likelihood of E_2 in the absence of E_1 (e.g. [9, 10]). Given that the stimulus and response function as the two events E_1 and E_2 , respectively, in sensorimotor learning, an associative account predicts that this form of learning should increase with $P(R/S)$ but decrease with $P(R/-S)$. This sensitivity to

contingency is captured by, among others, the Rescorla-Wagner model of associative learning [24].

Some years after publication of the ASL model, Keysers and Perrett [20] suggested that sensorimotor experience contributes to the development of the MNS via Hebbian learning. Hebb famously said that “Cells that fire together, wire together” and, more formally, “any two cells or systems of cells that are repeatedly active at the same time will tend to become ‘associated,’ so that activity in one facilitates activity in the other” [15, p. 70]. Thus, Keysers and Perrett’s Hebbian perspective implies that contiguity is sufficient for MNS development; that it does not also depend on contingency¹.

Hebbian learning is promiscuous; it connects any contiguously activated pair of cells or event representations. Consequently, Hebbian learning is at risk of supporting the establishment of internal connections that do not reflect reliable properties of the external world. For example, if one morning a person flexes her fingers in time to a piece of music, Hebbian learning could establish durable links between the finger movements and the music, even if she is just as likely to flex her fingers at other times, when the piece is not playing. By tracking the predictive relationship, or correlation, between events, contingency-based associative learning avoids this promiscuity problem. Therefore, the distinction between associative and Hebbian learning has important implications for the development of the MNS. In principle, associative learning based on sensorimotor experience is sufficient to explain why the MNS usually develops in a way that reflects real properties of the world; why it maps observed actions to the executed actions with which they systematically and reliably co-occur. In contrast, Hebbian learning based on sensorimotor experience is not sufficient to explain the observed properties of the MNS. To explain why the MNS tends to map observed actions to matching or ‘logically related’ executed actions, rather than to a large and semi-random set of executed actions – why the MNS is relatively free of ‘junk’ or ‘superstitious’ associations – it

is necessary to assume that Hebbian learning is guided by another process. For example, junk associations might be avoided if Hebbian learning was guided by the evolutionary process of canalisation; if individuals are born with nascent connections between sensory and motor neurons representing the same action, and sensorimotor experience merely facilitates the development of these connections, or if inborn mechanisms predispose individuals to seek sensorimotor experience from certain reliable sources [13]. Thus, associative learning could, but Hebbian learning could not, play a crucial, inductive role in the ontogeny of the MNS.

A recent study by Cook and colleagues [7] seems to provide evidence that sensorimotor experience modulates the operation of the MNS via associative rather than Hebbian learning. We will describe the background and procedure for this study in some detail because it is the focus of the computational modelling to be reported in this article. Cook et al. used ‘automatic imitation’ as an index of MNS responsivity. Automatic imitation is a stimulus-response compatibility effect in which the topographical features of task-irrelevant action stimuli facilitate similar, and interfere with dissimilar, responses [18]. (Topographic features of action relate to the way parts of the body move relative to one another, rather than to an external frame of reference.) Automatic imitation occurs even when it is contrary to task instructions, and incurs a financial cost [6]. Evidence that automatic imitation provides a valid index of MNS responsivity comes from research showing that repetitive transcranial magnetic stimulation (rTMS) of the inferior frontal gyrus – an area where mirror neurons have been found in monkeys – selectively disrupts automatic imitation [5].

The study by Cook and colleagues built on a previous experiment in which automatic imitation of opening and closing hand movements was measured before and after a period of sensorimotor training [19]. During training in this earlier study, the experimental group received novel, incompatible sensorimotor experience: whenever they observed an opening

hand stimulus, they made a closing hand response, and whenever they observed a closing hand stimulus, they made an opening hand response. The control group received familiar, compatible sensorimotor experience: they responded to opening hand stimuli with opening hand responses, and to closing hand stimuli with closing hand responses. In the pre- and post-tests, all participants completed a simple reaction time (RT) task in which they were required to make the same response (opening or closing) in every trial within a block, and to make this response as soon as they saw the stimulus hand begin to move. In compatible trials, the stimulus movement matched the pre-specified response (e.g. opening hand response in the presence of an opening hand stimulus), and in incompatible trials, the stimulus movement was the opposite of the pre-specified response (e.g. opening hand response in the presence of a closing hand stimulus). The results showed that the magnitude of the automatic imitation effect (calculated by subtracting RT on compatible trials from RT on incompatible trials) was smaller at post-test than at pre-test for the experimental group, but not for the control group. This outcome suggests that the incompatible sensorimotor training provoked – via either associative or Hebbian learning – the establishment of excitatory nonmatching connections (e.g. visual neurons activated by observation of hand opening linked in an excitatory way to motor neurons involved in hand closing) and/or the establishment of inhibitory matching connections (e.g. visual neurons activated by observation of hand opening linked in an inhibitory way to motor neurons involved in hand opening).

To investigate whether associative or Hebbian learning was responsible for the effect observed by Heyes et al., Cook and colleagues [7, Experiment 2] modified the previous experiment in three ways. First, they assessed automatic imitation only after sensorimotor training. Second, they gave all participants incompatible sensorimotor training. Third, the training received by the experimental group in the previous experiment, in which there was a perfect contingency between observation and execution of incompatible actions (contingent

group), was compared with ‘non-contingent’ and ‘signalled’ incompatible sensorimotor training. The number of contiguous sensorimotor pairings (432) was held constant across these three training groups. However, intermixed with these ‘paired trials’, in which an opening or closing action was executed in response to a numeric stimulus (1 or 2) and in the presence of the nonmatching action stimulus, the non-contingent and signalled groups each received an equal number of ‘unpaired’ trials (see Figure 1). In paired trials the numeric imperative stimulus was superimposed on an image of an opening or closing hand, whereas in unpaired trials it was superimposed on an image of a neutral hand posture. In the signalled group, unpaired trials were also distinguished from paired trials by contextual stimuli. Specifically, in unpaired trials, the neutral hand warning stimulus appeared in blue on a red background (rather than flesh colour on a black background), and was accompanied by a tone. Thus, contiguity was held constant across the three training groups – they experienced an equal number of pairings of action stimuli with nonmatching responses – while contingency and context were varied. The contingent group experienced a perfect, nonmatching sensorimotor contingency; the probability of a nonmatching response given an action stimulus was 1, whereas the probability of a nonmatching response in the absence of an action stimulus was 0. In contrast, the non-contingent and signalled groups experienced a zero, nonmatching sensorimotor contingency; the probability of a nonmatching response given an action stimulus was 1, but the probability of a nonmatching response in the absence of an action stimulus was also 1. Furthermore, in the signalled group, but not in the non-Contingent group, the trials that abolished the sensorimotor contingency – the unpaired trials – were presented in a distinctive context.

FIGURE 1 ABOUT HERE

After training in this experiment, Cook et al. [7] found a significantly larger automatic imitation effect in the non-contingent group than in either the contingent group or the

signalled group (see Figure 3, upper panel). In other words, non-contingent training left a larger residual automatic imitation effect, suggesting that the non-contingent group learned *less* as a result of incompatible sensorimotor training than the other two groups. The authors argued that this pattern of results implies that the effects of incompatible sensorimotor experience on automatic imitation were mediated by associative learning rather than Hebbian learning. However, their arguments were based on loose verbal specifications of the two theories, and informal inferences about the patterns of results that they would predict. Therefore, the current study examined more closely, using computational modelling, whether the findings reported by Cook and colleagues really support an associative over a Hebbian account of the way in which sensorimotor experience contributes to the development of the MNS.

Model Description

Basic Architecture

The study of Cook et al. [7] was simulated with a model consisting of ten units or nodes as shown in Figure 2. The model includes two sensory nodes which represent whether the observed hand movement on any trial is an opening or a closing movement; two imperative nodes which represent the numeric stimulus and hence whether the response required on the trial is an open response or a close response; four context nodes which represents context cues that may be present on the various trial types (the presence of each of the imperative cues, the warning stimulus on standard trials, and the warning stimulus on signalled trials) and that may therefore become associated with a response; and two motor nodes, corresponding to the open-hand and close-hand responses. The model operates according to the interactive activation principles of McClelland [21] as implemented in a previous model of imitative compatibility (or ‘automatic imitation’) effects [8]. Thus, each node has an activation level that varies between zero and one. Operation of the model is cyclic, with each cycle of the

model corresponding to a small interval of time. (In the simulations reported here, parameters are set such that one cycle corresponds to approximately one millisecond.) On each cycle, the activation a_i of each node i is calculated according to a simple difference equation:

$$a_i(t+1) = \rho \times a_i(t) + (1-\rho) \times \sigma(I_i(t)) \quad (1)$$

where t is the time, $I_i(t)$ is the net input to node i at time t , ρ is a parameter that controls the degree to which current activation persists from one cycle to the next, and $\sigma(\cdot)$ is the logistic or sigmoid function that maps all inputs to the range zero to one.

FIGURE 2 ABOUT HERE

Weighted connections between nodes allow nodes to excite or inhibit each other.

Following earlier work, we assume that associations between matching stimulus and response nodes are encoded as connections between sensory input nodes and motor response nodes acquired through learning in the course of everyday life and which mediate the propensity for automatic imitation that participants bring to the task. Moreover we assume that associations between imperative stimuli and corresponding responses are encoded as strong connections between imperative nodes and motor response nodes. These links, it is assumed, are set up by the subject in response to task instructions and maintained only for the duration of the experiment. Given this, the net input $I_j(t)$ to node j at time t is given by the sum of weighted excitation or inhibition to the node from other nodes (or from direct stimulation in the case of sensory, imperative, and context nodes), plus the node's bias (a parameter that sets the node's sensitivity to excitatory input), plus normally distributed noise. That is:

$$I_j(t) = \sum_i (w_{ji} \times a_i(t-1)) + E_j + \beta_j + N(0, \eta^2) \quad (2)$$

where w_{ji} is the strength of the connection from node i to node j , E_j is any direct stimulation applied to the node, β_j is the bias on node j , and η is the standard deviation of noise added on each processing cycle.

When a stimulus appears, direct stimulation is applied to corresponding sensory nodes, an imperative node, and any relevant context nodes. This causes the activation of those nodes to increase (by equation 1), resulting in increased input to motor nodes (by equation 2). A response is assumed to be generated when the activation of a response node exceeds a threshold, θ , which following earlier work [8] is set to 0.8 in the simulations reported below. The number of processing cycles between stimulus presentation and response generation is assumed to be proportional to the time taken by a subject to produce the analogous response to the analogous stimulus.

Learning Rules

The architecture as described above is capable of simulating imitative stimulus-response compatibility effects (cf. [8]). In order to extend it to the Cook et al. study it is necessary to supplement it with a learning rule that specifies how the strengths of S-R associations are modified in response to events such as the co-occurrence of a stimulus and a response. Generalising from the introduction, we consider three possible learning rules:

1. Standard Hebbian learning: The association between a stimulus S_i and a response R_j is strengthened if the stimulus and response co-occur. If S_i occurs in the absence of R_j (or the presence of another response), or R_j occurs in the absence of S_i (or the presence of another stimulus), then the $S_i - R_j$ association is unchanged. In symbols:

$$\Delta w_{ji} = \begin{cases} 1 & \text{(if } S_i \text{ and } R_j \text{ are both present)} \\ 0 & \text{(otherwise)} \end{cases}$$

Hebbian learning is adjusted to prevent weights from growing without bound by incorporating an asymptote, α :

$$\Delta w_{ji} = \begin{cases} (\alpha - w_{ji}) / w_{ji} & \text{(if } S_i \text{ and } R_j \text{ are both present)} \\ 0 & \text{(otherwise)} \end{cases} \quad (3a)$$

2. Quasi-Hebbian learning: This is basically Hebbian learning with an adjustment for non-contingent stimulus-response occurrences. As in standard Hebbian learning, the

association between a stimulus S_i and a response R_j is strengthened if the stimulus and response co-occur. However, if R_j occurs in the absence of S_i , or S_i occurs in the absence of R_j , then the $S_i - R_j$ association is weakened. In symbols:

$$\Delta w_{ji} = \begin{cases} 1 & \text{(if } S_i \text{ and } R_j \text{ are both present)} \\ -1 & \text{(if } S_i \text{ is present but } R_j \text{ absent or vice versa)} \\ 0 & \text{(otherwise)} \end{cases}$$

As with Hebbian learning, this is adjusted to incorporate an asymptote, α :

$$\Delta w_{ji} = \begin{cases} (\alpha - w_{ji}) / w_{ji} & \text{(if } S_i \text{ and } R_j \text{ are both present)} \\ -(\alpha - w_{ji}) / w_{ji} & \text{(if } S_i \text{ is present but } R_j \text{ absent or vice versa)} \\ 0 & \text{(otherwise)} \end{cases} \quad (3b)$$

We consider quasi-Hebbian learning because it allows for a form of extinction, whereby absence of a predicted response results in a reduction in the strength of association between the stimulus and the prediction. As we will see, this addresses some of the limitations of purely Hebbian learning in accounting for the effects observed by Cook et al. It does this while remaining close in spirit to Hebb's original formulation.

3. Associative (or Rescorla-Wagner) learning: In this case, changes in association strength are proportional to the error (ϵ) between a target value (τ) of a response node and the actual value given the stimulus. Where there are multiple simultaneous stimuli, this error is attributed to, or distributed over, the various sources of activation on any trial in proportion to the strength of those activation sources. In symbols:

$$\begin{aligned} \epsilon_j &= \tau_j - \sum_i (w_{ji} \times a_i) \\ \Delta w_{ji} &= \epsilon_j \times a_i \end{aligned} \quad (3c)$$

On a trial in which R_j is present, ϵ_j is set to the input necessary to drive the response node to its maximum. Critically, if the target is predicted by the input, error will be zero and no learning will occur, even if a novel stimulus is present on that trial. Equally critically, the target value may be over-predicted, leading to weakening of the association between S_i and R_j ,

even when S_i and R_j are both present. Associations between any stimuli present on a trial and the response R_j will also be weakened if the response is predicted but absent (i.e., if $\sum_i (w_{ji} \times a_i) > 0$ but $\tau_j = 0$). This will also result in ε_j being negative and any associations between stimuli that are present on that trial will be weakened.

For each of these learning rules, we assume that the context node associated with one of the imperative stimuli is active with strength C_{imp} on every trial of training (depending on which imperative stimulus is presented on that trial) and acts like just another stimulus, so that associations between the context node and the response nodes may be learned. We also assume that the context node associated with the warning stimulus is active with strength C_1 or C_2 , depending on whether the trial is a standard training trial (C_1) or a signalled trial (C_2). Moreover for the two forms of Hebbian learning we assume that the change in any association strength on each trial is scaled by the magnitude of the difference between the association's current strength and an asymptotic value (α) so that association strengths are limited to that asymptotic value.

The addition of learning introduces a further five parameters over and above the non-learning model of Cooper et al. [8]:²

- λ , the learning rate (a positive real number, ranging from 0.0002 to 0.0060 in the simulations reported below) that scales the weight adjustment (Δw_{ji}) made on each learning trial³
- α , the weight asymptote for Hebbian and quasi-Hebbian learning or τ , the target input required to drive a node to its maximum for Rescorla-Wagner learning (both positive real numbers, ranging from 2 to 20 in the simulations reported below)
- C_{imp} , C_1 and C_2 , the strength of the various context nodes, relative to the strength of the irrelevant sensory stimuli (positive real numbers, fixed at 1 for C_{imp} and C_1 and ranging from 4 to 12 for C_2 in the simulations reported below)

Simulation Study 1

The purpose of simulation study 1 was to determine, for each learning rule, whether the model could account for the Cook et al. data, i.e., whether the model could account for a greater residual automatic imitation effect following non-contingent training than following contingent or signalled training.

Method

The model contains three kinds of parameters: those that concern the general activation dynamics (e.g., the bias on all nodes and the persistence of activation over time), those that specify the initial strengths of associations, and those related to learning (e.g., the learning rate and asymptotic association strength). In the simulations reported here, all but one of the parameters associated with general activation dynamics were held at the values used in previous work [8].⁴ The initial strengths of associations were also set based on this previous work to +10 for associations between imperative nodes and corresponding response nodes and +4 for associations between sensory nodes and compatible response nodes. The former reflects the deliberate activation of motor response nodes following an imperative stimulus, while the latter reflects the automatic activation of compatible response nodes following presentation of a sensory stimulus. Following previous work with these values we assume that one processing cycle of the model corresponds to approximately one millisecond of subject time.

For each of the three learning rules, a series of simulations was conducted varying three learning parameters: the learning rate, the weight asymptote or equivalent (i.e., α or τ), and the relative strength of the signalled context. In each case, a “virtual subject” was simulated by replicating the training and testing experienced by a real subject from the Cook et al. experiment. Thus, the model was a) initialised, then b) trained for 6 blocks with one of the three training sets (72 trials per block for the contingent condition, 144 trials per block for

the non-contingent and signalled conditions), and finally c) tested for 60 instances of each compatible and incompatible stimulus-response pairing. For each learning rule this procedure was repeated for all three training sets. The compatibility effects following each training set were then calculated yielding three values directly comparable to the compatibility effects in the subject data (20.8 msec, 33.6 msec, and 18.6 msec for contingent, non-contingent and signalled training, respectively). Finally, the fit of the model to the data was calculated as the root mean square (RMS) difference between the three pairs of values.

The fit of the model to the data is a function of the learning parameters. Therefore for each learning rule this procedure was repeated for all values within a three dimensional grid with the learning rate ranging from 0.0002 to 0.0060 in steps of 0.0002, the weight asymptote or equivalent ranging from 2 to 20 in steps of 2, and the relative strength of the signalling context ranging from 4 to 12 in steps of 2. Furthermore the model's behaviour is non-deterministic, because the order of training trials is randomised prior to each training phase and initial node activations are randomised prior to each testing trial. Therefore the whole procedure was repeated 10 times for each point in the parameter space, with the RMS fit averaged over these 10 replications. This procedure yielded a total of $30 \times 10 \times 5 \times 10 = 15000$ simulations per learning rule. The Express software [26] was used to manage the exploration of the parameter space and to collate results.

Results and Discussion

Table 1 shows the parameter values and mean RMS score for the best fitting model for each rule. By varying the three learning parameters it is possible to obtain apparently reasonable fits (within 10 msec RMS) for all learning rules, though it appears from the table that the best fit is obtained with the Rescorla-Wagner learning rule. However, these apparently good fits are in some ways misleading. Figure 3 shows the compatibility effects obtained using each of the parameter settings from Table 1. It is clear from the figure that fit of the model

using Hebbian learning is in fact poor. The fact that reaction time (RT) of the model expressed in cycles is approximately 30 less than RT of subjects expressed in milliseconds is not of great concern. Even with processing at approximately 1 millisecond per cycle, it is reasonable to assume that additional pre or post decision processes, not modelled, would bring the model into line with mean human RT on the task. Critically, however, Hebbian learning fails to reproduce the appropriate compatibility effects – the predicted compatibility effect is greater with contingent training or with signalled training than with non-contingent training. This is precisely opposite to the result obtained in the Cook et al. study.

FIGURE 3 ABOUT HERE

TABLE 1 ABOUT HERE

Quasi-Hebbian learning fares better in one respect but worse in another. Quasi-Hebbian learning can capture the fact that the compatibility effect is smaller with contingent training than with non-contingent training – and this justifies our consideration of it as an alternative to standard Hebbian learning – but it suggests that RT should be slower with non-contingent training than with contingent training (even on compatible trials), and that signalled training should yield a similar compatibility effect to non-contingent training. Both of these effects are contradicted by the subject data.

None of the difficulties exhibited by the model using Hebbian or quasi-Hebbian learning are shown with Rescorla-Wagner learning. Here, as in the human data, the compatibility effect is greater with the non-contingent training set than with either of the other training sets. Moreover, the compatibility effect is similar with contingent training and signalled training. That RT (in cycles) is generally slower than in the subject data (in seconds) is also not a major concern – this may be addressed by assuming that the strength of the imperative stimulus (the numerals 1 or 2) is slightly greater in this study than in the study of

Catmur and Heyes [3] (an orange or purple dot), on which the earlier model (and the strength of imperative to motor associations) was based.

Simulation Study 2

One difficulty with simulation 1 is that in finding the best fit in quantitative terms of the model to the data for each learning rule, we may have failed to note the potential for Hebbian or quasi-Hebbian learning to account, in principle, for the qualitative effects. A second difficulty is that the model with Rescorla-Wagner learning may be so powerful, and have so many free parameters, that it could in principle account for any pattern of results. If this were the case it would significantly reduce the explanatory power of the model [25]. Simulation study 2 therefore adopts a method based on *parameter space partitioning* [22] to explore the potential ability of each learning rule to account for the critical qualitative effects.

Parameter space partitioning aims to divide the parameter space of a computational or mathematical model into regions corresponding to qualitatively distinct behaviours. Each point in the parameter space is categorised according to the pattern of effects present in the model's behaviour at that point. Within the context of the Cook et al. study, there are three effects of concern. First, the compatibility effect is significantly greater with non-contingent training than with contingent training. Second, the compatibility effect is significantly greater with non-contingent training than with signalled training. Third, the compatibility effects with contingent training and with signalled training are not significantly different. (Strictly speaking the last of these is the absence of an effect, but it would be worrying if the model were to predict such an effect without it having been found in the data from human subjects.) Any point in the model's parameter space might yield any combination of these effects. If, across the parameter space, all combinations are possible then the model is of little explanatory value – it could account for any pattern of data. If on the other hand only the

observed combination is possible then the model has strong explanatory value – any settings of its parameters would yield the observed qualitative pattern of results.

Method

For simulation study 1 it was sufficient to simulate the behaviour of a single subject at different points in the model's parameter space. Simulation study 2 requires the simulation of a group of subjects (including between-subject differences), so that standard statistical analyses may be performed on the resulting group data in order to determine whether type of training has a statistically significant effect on the compatibility effect. To do this, the strength of the short-term associations between imperative nodes and motor response nodes was sampled from a normal distribution with mean 10 and standard deviation 1. This association can be seen as reflecting individual subjects' motivation or commitment to the task.

Increasing the value increases the speed of response to all stimuli, while decreasing it leads to slower response times.

For each of the three learning rules, a series of simulations was conducted varying the learning rate and the weight asymptote or equivalent. To maintain tractability the relative strength of the signalled context, C_2 , was fixed at 10.0. (Inspection of the results from simulation study 1 suggests that this does not substantially compromise model fit.) In each case, each individual simulation consisted of 12 virtual subjects (mirroring the 12 subjects per condition in the Cook et al. study), with the procedure for each virtual subject being as in simulation study 1. All other parameters beyond those being varied were fixed at the values used in simulation study 1 (with the exception of sampling the strength of imperative to motor response associations as described in the previous paragraph).

Parameter space partitioning focuses on qualitative effects. Therefore, rather than calculating quantitative compatibility effects, for each simulation we calculated whether each of the three effects of interest was statistically significant (given the relevant null hypothesis).

Thus, between-subjects t-tests were performed comparing a) the compatibility effect with contingent training versus non-contingent training, b) signalled training versus non-contingent training, and c) contingent training versus signalled training. Results of each t-test were scored as +1, 0 or -1 for each simulation, depending on whether the effect was significant and positive ($p < 0.05$, one-tailed), non-significant, or significant and negative ($p < 0.05$, one-tailed), respectively. Note that for the model to replicate the qualitative pattern in the observed data it must score -1, -1 and 0, respectively, on these three measures.

The simulation was run for all pairs of parameter values within a two dimensional grid with the learning rate ranging from 0.0000 to 0.0060 in steps of 0.0002 and the weight asymptote or equivalent ranging from 2 to 22 in steps of 2. Since the model's behaviour is non-deterministic (with the variable strength of associations from imperative to motor response nodes adding to the variability in response times), the whole procedure was repeated 100 times for each point in the parameter space, with the mean score for each statistical test calculated over the 100 replications. This procedure yielded a total of $31 \times 11 \times 100 = 31400$ simulations per learning rule. As before, the Express software [26] was used to manage the exploration of the parameter space and to collate results.

Results and Discussion

Results of simulation study 2 are summarised in Figure 4. With Hebbian learning, most pairs of parameter settings result in the compatibility effect being significantly smaller with non-contingent training than with contingent training (blue region in Figure 4, upper left plot). This is consistent with simulation study 1 but opposite to what was found by Cook et al. [7]. Only when the learning rate is very low is there no significant effect of training type (green region in Figure 4, upper left plot), but at this low rate of learning there is minimal learning in either condition. Similar results hold for the comparison of compatibility effects with non-contingent versus signalled training (Figure 4, upper centre plot). Hebbian learning also

generally predicts that signalled training will result in a compatibility effect that is significantly smaller than that arising from contingent training (blue region in Figure 4, upper right plot). Thus, the inability of Hebbian learning in simulation study 1 to replicate the effects in the subject data is not due to selection of sub-optimal parameter values: across the majority of the parameter space Hebbian learning predicts that non-contingent training will result in a smaller compatibility effect than signalled training, which in turn will result in a smaller compatibility effect than contingent training.

FIGURE 4 ABOUT HERE

The central panel of Figure 4 depicts the results with quasi-Hebbian learning. The addition of an anti-Hebbian term to the learning equation allows the model to account for the effect of non-contingent versus contingent training on the compatibility effect – consistent with the subject data, this version of the model predicts that the compatibility effect will be greater with non-contingent than contingent training. This is true for the vast majority of the parameter space (red region, middle left plot). However, quasi-Hebbian learning fails to account for the effect of signalled training on the compatibility effect. The learning rule predicts that non-contingent and signalled training will have similar effects (green region, middle centre plot), contrary to the subject data. Quasi-Hebbian learning also predicts that signalled training will generally result in a greater compatibility effect than contingent training (red region, middle right plot). Again this is contrary to what was found by Cook et al. [7].

Results from Rescorla-Wagner learning are shown in the lower panel. Here, a large region of parameter space yields models that result in greater compatibility effects with non-contingent than with contingent training (red region, lower left plot). A subset of this region yields models that result in greater compatibility effects with non-contingent than with signalled training (red region, lower centre plot). An overlapping region of parameter space

yields models that result in similar compatibility effects with signalled and contingent training (green region, lower right plot). The intersection of these three regions corresponds to parameter settings that yield patterns of effects that are qualitatively equivalent to those found by Cook et al. [7].

FIGURE 5 ABOUT HERE

Figure 5 summarises the overlaps between the three qualitative fits for each learning rule. Red regions correspond to parameter settings where all three effects are, qualitatively, as found in the subject data. As should be clear from the above description of results, if Hebbian or quasi-Hebbian learning is assumed, then it is not possible to select parameters where the model replicates the observed behavioural effects. This is only possible with Rescorla-Wagner learning. However, Rescorla-Wagner learning does not guarantee that the model will produce the observed behaviour: this occurs only in a crescent-shaped region of parameter space. If the target value (τ in Equation 3c) is relatively high, then the learning rate must be relatively low, but at low values of τ higher learning rates are required to reproduce the qualitative effects found in the human data. In the region to the top-right of this crescent, where the model fails to reproduce the qualitative pattern in the human data, it can be seen from Figure 4 that this is because in this region of the parameter space the model predicts that the compatibility effect with signalled training should be greater than that with contingent training. We return to why this is so in the General Discussion.

General Discussion

The simulation results demonstrate firstly that the sensorimotor compatibility effects such as those found by Heyes et al. [19], Cook et al. [7] and Catmur and Heyes [3] can be accounted for in terms of weighted associations between sensory and motor units. The work therefore provides additional support for the model underlying the previous simulation of Cooper et al. [8]. More critically in the context of discussions of the MNS, the simulation results

demonstrate that associative learning can, and Hebbian learning cannot, account for the effects of different types of training on the automatic imitation effect as reported by Cook et al.

The simulation results generally follow the logic proposed by Cook et al. [7] in their discussion of the putative effects of contingent, non-contingent and signalled training. However, the informal verbal reasoning of Cook et al. considers only how stimulus-response associations are altered depending upon the probability of a response given that a stimulus is present, $P(R/S)$, or that a stimulus is absent, $P(R/\sim S)$. It does not consider how trials on which an expected response is absent affect association strengths. That is, it does not consider the effect of $P(\sim R/S)$. The computational model addresses this potential gap in the reasoning. (See the Supplementary Materials for detailed discussion of how these events affect the associative weights under the different learning rules.)

The other notable feature of the simulation results is the tendency for the model with Rescorla-Wagner learning, high learning rate (e.g., $\lambda = 0.005$) and high target (e.g. $\tau = 20.0$) to predict a larger residual compatibility effect in the signalled condition than in the contingent condition (cf. the red region in Figure 4, lower right panel). If parameter space partitioning had not indicated that there are many potential outcomes of the Cook et al. experiment that Rescorla-Wagner learning could *not* accommodate, this feature of the simulation results might be regarded as a weakness of the Rescorla-Wagner model. However, since parameter space partitioning demonstrated that the Rescorla-Wagner model is eminently falsifiable, this feature constitutes a novel prediction – that there will be specific parameter values at which signalled training yields a larger residual compatibility effect than contingent training – to be tested in future experiments.

In conclusion, the simulation studies reported in this article show that, of the learning algorithms considered, only Rescorla-Wagner learning can account for the pattern of results

found in the study of Cook et al. Therefore, not only do they provide negative evidence for the hypothesis that sensorimotor experience contributes to the development of the MNS via Hebbian (or quasi-Hebbian) learning, they also provide positive evidence for the hypothesis that sensorimotor experience contributes to the development of the MNS via associative (i.e., prediction-error) learning. They do this by showing both that Rescorla-Wagner learning can account for the data, and, using parameter space partitioning, that this result was very far from inevitable. More broadly, our findings support the associative sequence learning model (ASL), which accords a crucial, inductive role to sensorimotor experience in the development of the MNS [16, 17]. Hebbian learning would require guidance, or canalisation, to yield the documented properties of the mature MNS, but associative learning is potentially sufficient to build a MNS that maps observed actions to the executed actions with which they systematically and reliably co-occur.

References

- [1] Calvo-Merino, B., Grezes, J., Glaser, D. E., Passingham, R. E., & Haggard, P. (2006). Seeing or doing? Influence of visual and motor familiarity in action observation. *Current Biology*, *16*, 1905-1910.
- [2] Catmur, C. (in press). Sensorimotor learning and the ontogeny of the mirror neuron system. *Neuroscience Letters*.
- [3] Catmur, C., & Heyes, C. M. (2011). Time course analyses confirm independence of automatic imitation and spatial compatibility effects. *Journal of Experimental Psychology: Human Perception and Performance*, *37*, 409-421. DOI: 10.1037/a0019325
- [4] Catmur, C., Mars, R., Rushworth, M., & Heyes, C. M. (2011). Making mirrors: premotor cortex stimulation enhances mirror and counter-mirror motor facilitation. *Journal of Cognitive Neuroscience*, *23*, 2352–2362.
- [5] Catmur, C., Walsh, V., & Heyes, C. M. (2009). The role of experience in the development of imitation and the mirror system. *Philosophical Transactions of the Royal Society B*, *364*, 2369 - 2380.
- [6] Cook, R., Bird, G., Lunser, G., Huck, S. & Heyes, C. M. (2011). Automatic imitation in a strategic context: Players of Rock-Paper-Scissors imitate opponents' gestures. *Proceedings of the Royal Society of London: B*. doi: 10.1098/rspb.2011.1024
- [7] Cook, R., Press, C., Dickinson, A., & Heyes, C. M. (2010). Acquisition of automatic imitation is sensitive to sensorimotor contingency. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(4), 840-852.
- [8] Cooper, R. P., Catmur, C., & Heyes, C. M. (in press). Are automatic imitation and spatial compatibility mediated by different processes? *Cognitive Science*.

- [9] Dickinson, A., & Charnock, D. J. (1985). Contingency effects with maintained instrumental reinforcement. *Quarterly Journal of Experimental Psychology*, 37B, 197-416.
- [10] Elsner, B., & Hommel, B. (2004). Contiguity and contingency in action effect learning. *Psychological Research*, 68, 138-154.
- [11] Ferrari, P. F., Rozzi, S., & Fogassi, L. (2005). Mirror neuron responding to observation of actions made with tools in monkey ventral premotor cortex. *Journal of Cognitive Neuroscience* 17, 212-226.
- [12] Gillmeister, H., Catmur, C., Liepelt, R., Brass, M., & Heyes, C. M. (2008). Experience-based priming of body parts: A study of imitation and the mirror system. *Brain Research*, 1217, 157-170.
- [13] Del Giudice, M., Manera, V., & Keyesers, C. (2009). Programmed to learn? The ontogeny of mirror neurons. *Developmental Science*, 12, 350-363.
- [14] Haslinger, B. et al., 2005. Transmodal sensorimotor networks during action observation in professional pianists. *Journal of Cognitive Neuroscience* 17, 282-293.
- [15] Hebb, D. (1949). *The Organization of Behaviour*. New York: Wiley.
- [16] Heyes, C. M. (2001). Causes and consequences of imitation. *Trends in Cognitive Sciences* 5, 245-261.
- [17] Heyes, C. M. (2010). Where do mirror neurons come from? *Neuroscience and Biobehavioral Reviews*, 34, 575-583.
- [18] Heyes, C. M. (2011). Automatic imitation. *Psychological Bulletin*, 137, 463-483.
- [19] Heyes, C. M., Bird, G., Johnson, H., & Haggard, P. (2005). Experience modulates automatic imitation. *Cognitive Brain Research*, 22, 233-240.
- [20] Keyesers, C., & Perrett, D. I. (2004). Demystifying social cognition: a Hebbian perspective. *Trends in Cognitive Sciences* 8, 501-507.

- [21] McClelland, J. L. (1993). The GRAIN model: A framework for modeling the dynamics of information processing. In D. E. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience* (pp. 655–688). Hillsdale, NJ: Erlbaum.
- [22] Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review*, *113*, 57-83.
- [23] Press, C., Gillmeister, H., & Heyes, C. M. (2007). Sensorimotor experience enhances automatic imitation of robotic action. *Proceedings of the Royal Society of London, Biological Sciences* *274*, 2509-2514.
- [24] Rescorla, R., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical Conditioning II: current research and theory* (pp.64-99). New York: Appleton-Century-Crofts.
- [25] Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, *107*(2), 358-367.
- [26] Yule, P., & Cooper, R. P. (2003): Express: A web-based technology to support human and computational experimentation. *Behavior Research Methods, Instruments, and Computers*, *35*, 605–613.

Footnotes

¹ In some contexts, Hebbian learning could be regarded as a species of associative learning. The terms ‘associative’ and ‘Hebbian’ are used contrastively here to highlight the differences between the two types of learning (based on contiguity and contingency, or on contiguity alone) and between the ASL (e.g. [1]) and Hebbian [20] accounts of the development of the MNS.

² Full details of all model parameters are given in the supplementary materials. The complete model, which is written in the C programming language, is available for download from <http://www.ccnl.bbk.ac.uk/models.html>

³ Thus, if $w_{ij}(t)$ is the weight of the association to node j from node i at time t , then for each learning trial $w_{ij}(t+1) = w_{ij}(t) + \lambda \Delta w_{ji}$

⁴ The one parameter that was not held at its value from previous work was the habituation threshold of sensory nodes. In the previously reported work, sensory nodes were assumed to habituate to an input and decay once their activation exceeded a value of 0.80. In the current work this threshold was increased to 0.90 in order to yield effect sizes similar to that seen in the empirical work of Cook et al. [7, Experiment 2]. See the Supplementary Materials for details of all model parameters.

Figures:

Figure 1

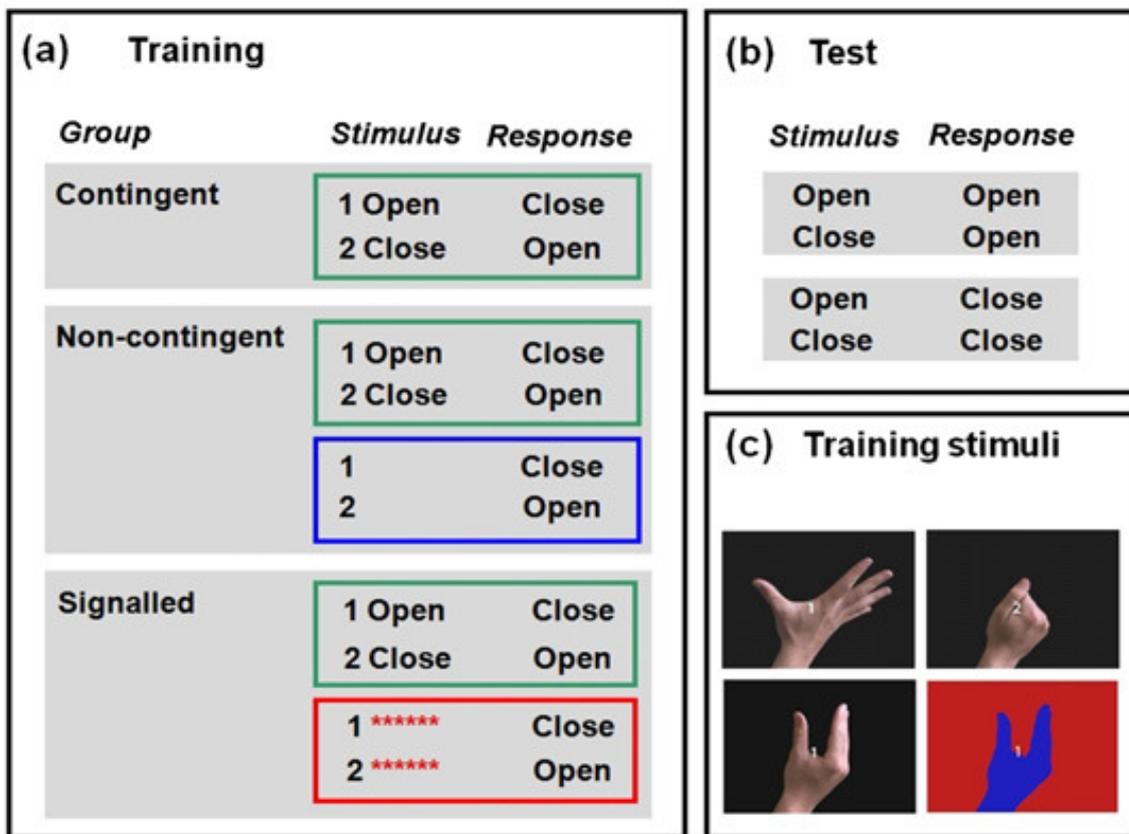


Figure 1: Summary of the experiments reported by Cook et al. [7]: (a) participants were allocated to one of three counter mirror training groups (contingent, non-contingent or signalled). All groups received the same number of paired trials (green boxes) where the execution of a response (e.g. open hand) was paired with the observation of the counter-mirror stimulus (close hand). In the non-contingent group the sensorimotor contingency was degraded through the addition of unpaired trials (blue box) where participants executed responses, not in the presence of the counter-mirror stimulus, but while observing a neutral hand warning stimulus. In the signalled group, the contingency was degraded through the addition of signalled trials (red box), where responses were made in the presence of a differentiated neutral hand warning stimulus. Participants completed six blocks of counter-mirror training trials spread evenly over a two day period, (b) following training, all participants completed an identical test procedure to measure the size of their residual automatic imitation effects. Participants were required to execute open- and close-hand responses to the onset of open- (compatible) and close-hand (incompatible) stimuli. Automatic imitation effects were estimated by subtracting mean RTs on compatible trials from mean RTs on incompatible trials. Smaller residual automatic imitation effects were seen following contingent training, indicating that this was the most effective schedule and (c) the stimuli used during training. Top row: stimuli used on paired trials. Bottom left: the warning stimulus presented on unpaired trials in the non-contingent training schedule. Bottom right: the differentiated warning stimulus presented on signalled trials.

Figure 2

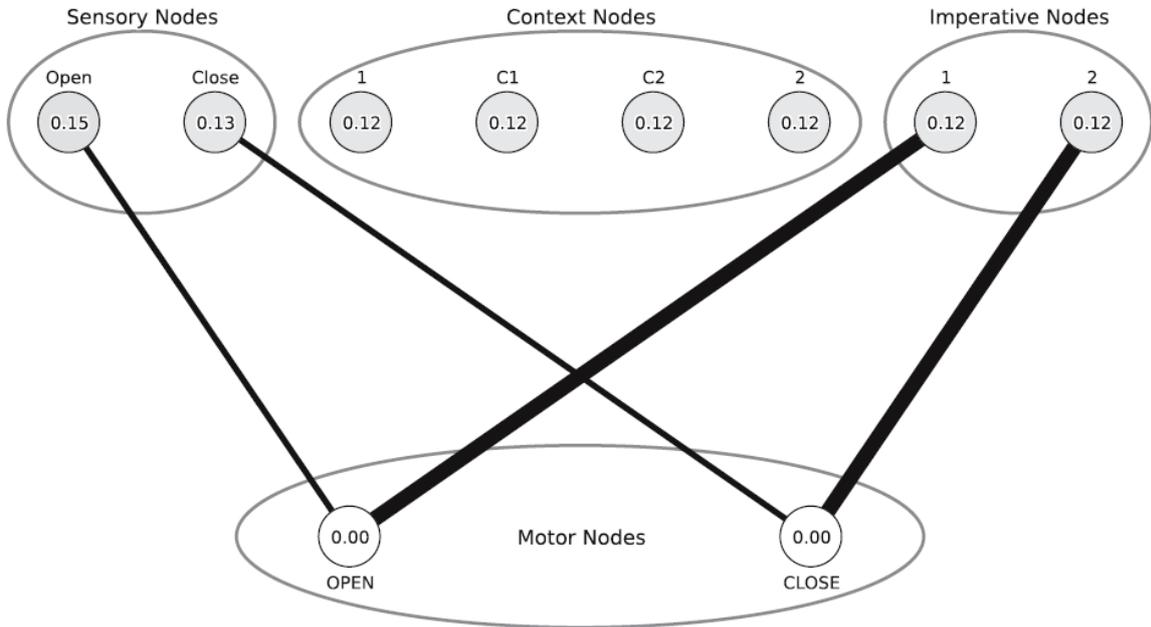


Figure 2: Architecture of the model, showing associations between nodes prior to learning.

Figure 3

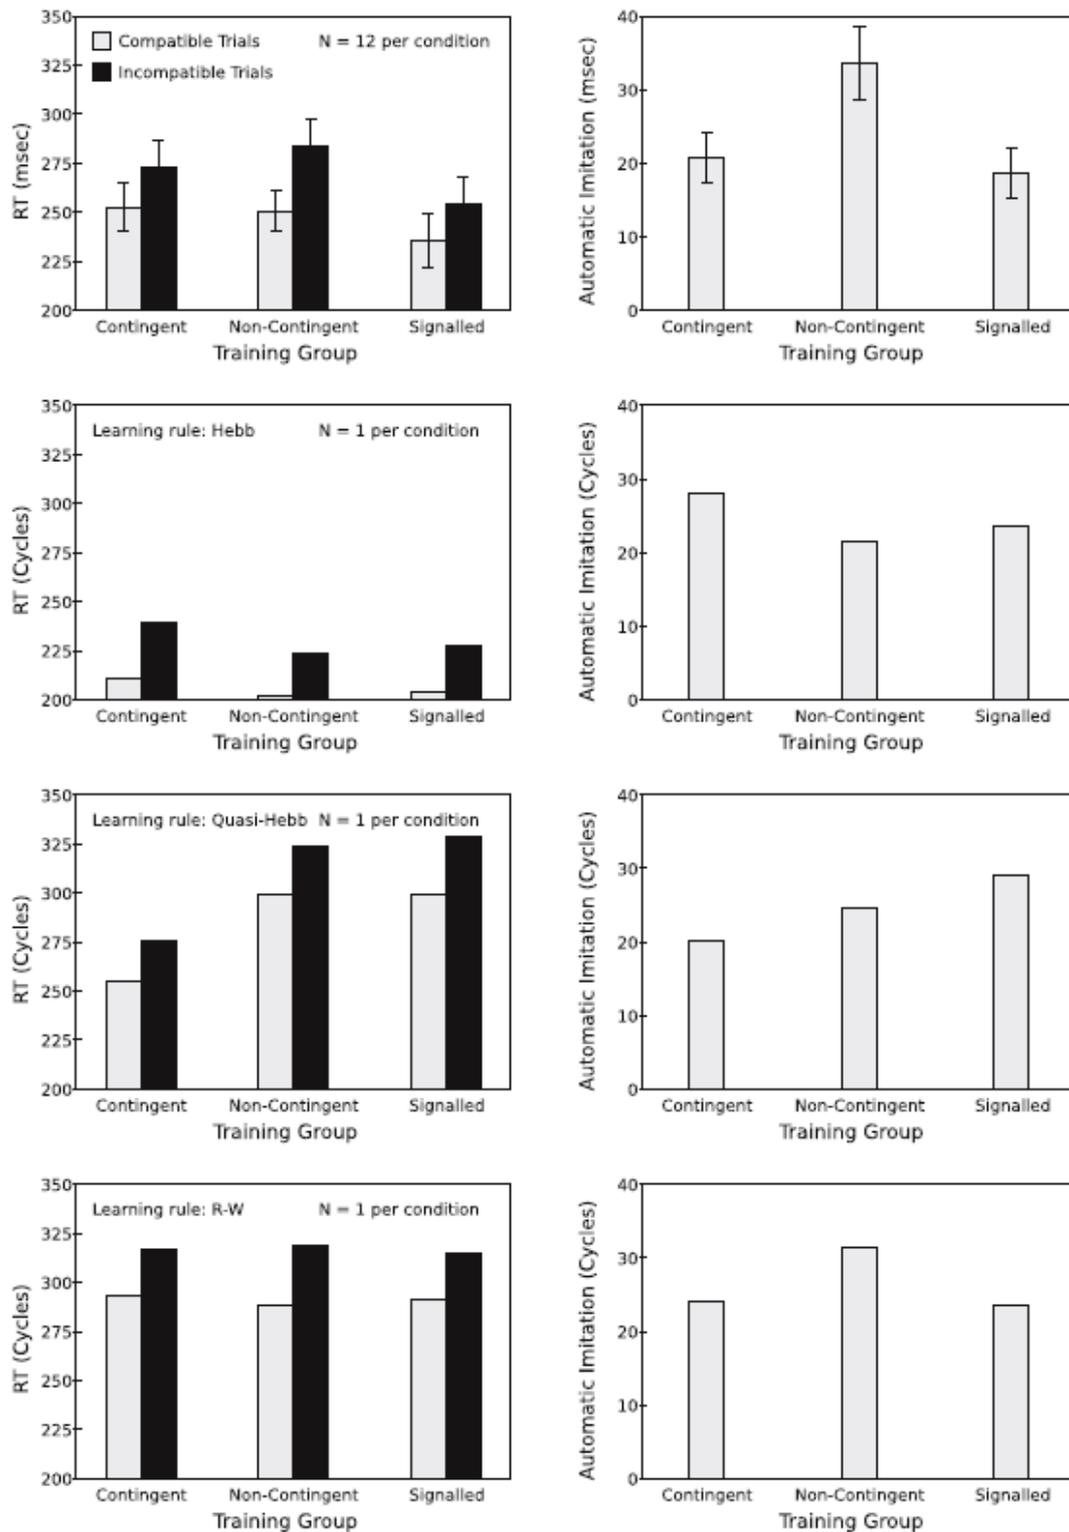


Figure 3: Compatibility effects for subject and simulated data. The top-most panel shows the subject data of Cook et al. [7, Experiment 2], with the reaction time (RT) data for compatible and incompatible trials for each training group on the left, and the compatibility effects (i.e., different between incompatible and compatible reaction time) on the right. Error bars show standard error of the main. The lower panels show best fits for each learning (Hebb, Quasi-Hebb and Rescorla-Wagner respectively). Only in the case of Rescorla-Wagner learning are the critical empirical effects (similar compatibility effects with contingent and signalled training, but a significantly greater compatibility effect with non-contingent training) reproduced.

Figure 4

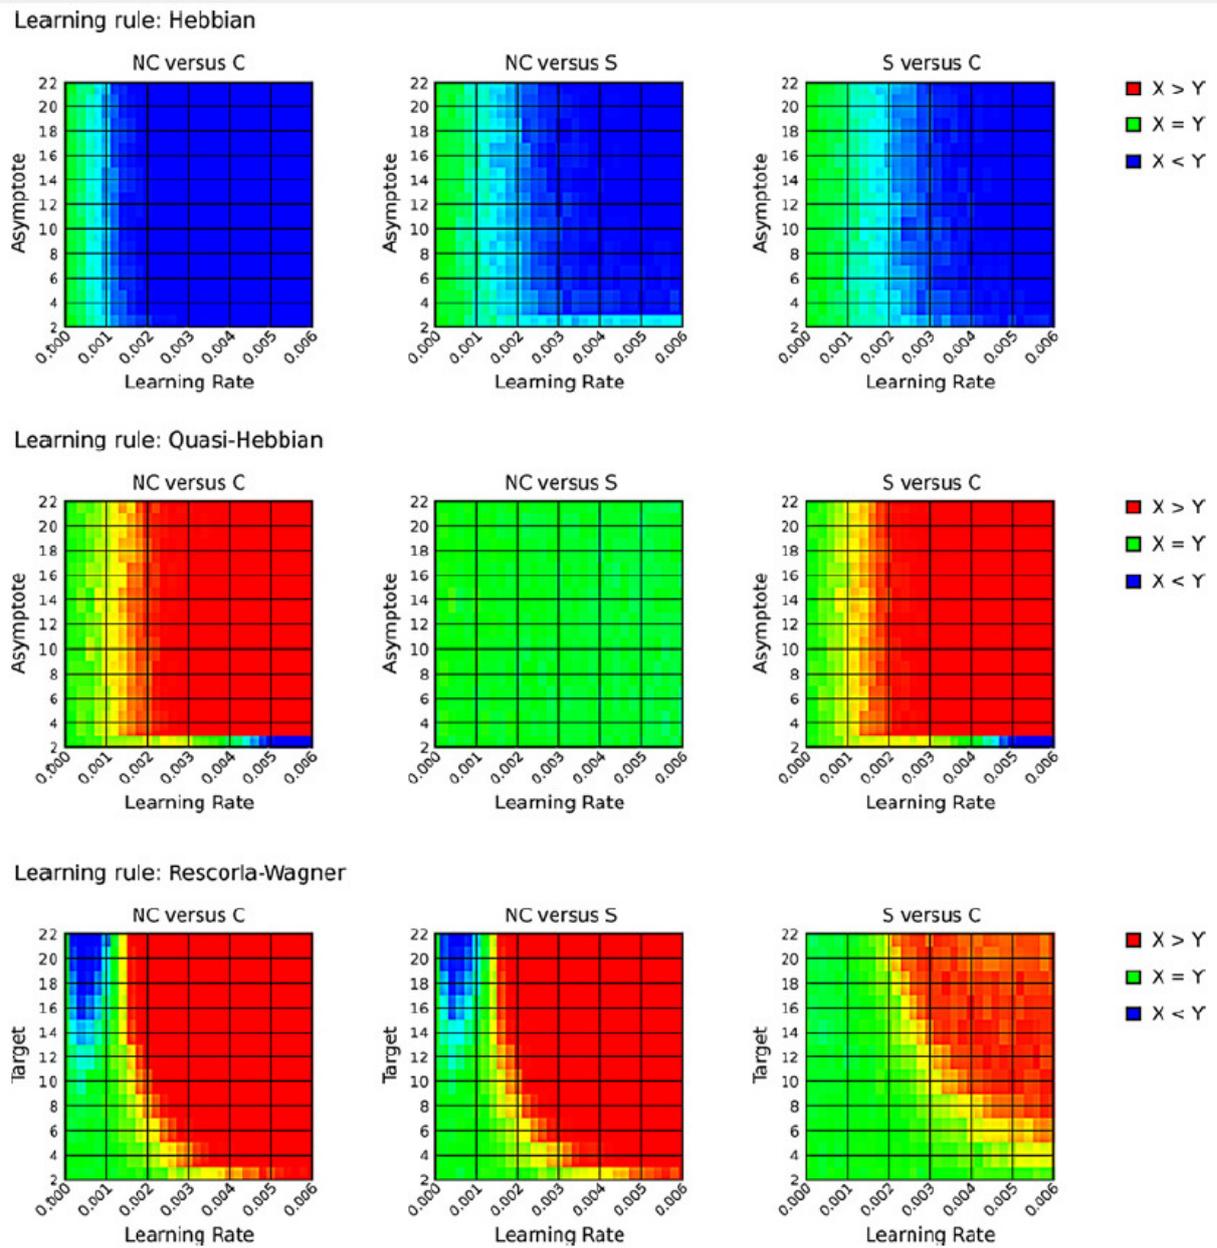


Figure 4: Qualitative effects exhibited by the model with Hebbian (upper panel), Quasi-Hebbian (middle panel), and Rescorla-Wagner (lower panel) learning at different points in the parameter space. In each case the left figure shows where the compatibility effect differs with non-contingent training versus contingent training. In the blue region, the effect is significantly smaller with non-contingent than with contingent training, while in the red region the effect is significantly greater with non-contingent than with contingent training. Recall that the latter was observed in the Cook et al. study. The central column shows the equivalent comparison for non-contingent training versus signalled training. Again, in the subject data the compatibility effect was significantly greater with non-contingent than with signalled training (corresponding again to the red region). The right column shows the comparison between signalled and contingent training. In this case no effect was found in the subject data (corresponding to the green region).

Figure 5

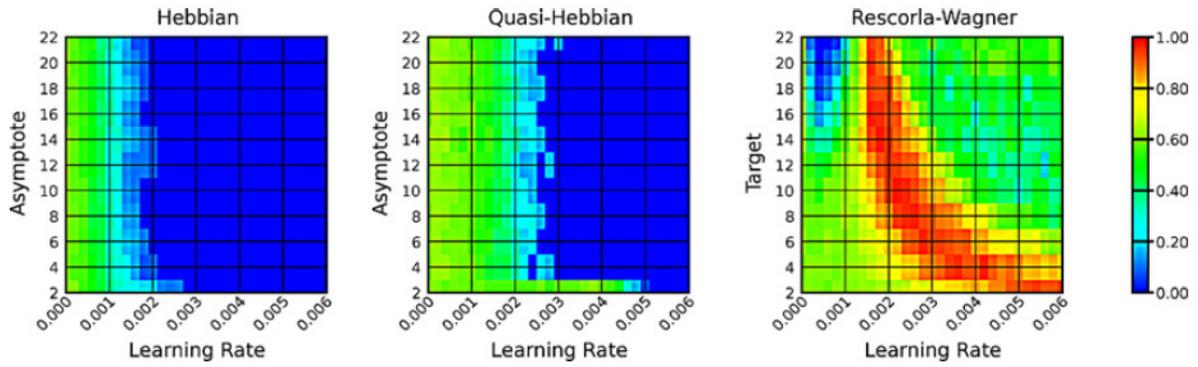


Figure 5: Fit of the model with different learning rules to the qualitative effects in the subject data of Cook et al. [7, Experiment 2]. These plots effectively show the ability of the model to simultaneously fit all three effects considered in Figure 4. Red regions correspond to parameter settings where the model replicates all effects found in the subject data. They correspond to the overlap between red regions in Figure 4 left, red regions in Figure 4 centre, and green regions in Figure 4 right, for each learning rule.